arXiv:1411.0031v2 [stat.ME] 29 Nov 2014

Likelihood-Based Inference for Discretely Observed Birth-Death-Shift Processes, with Applications to Evolution of Mobile Genetic Elements

By Jason Xu*, Peter Guttorp*, Midori Kato-Maeda†, Vladimir Minin*,‡

*Department of Statistics, University of Washington

‡Department of Biology, University of Washington

†School of Medicine, University of California, San Francisco

Abstract

Continuous-time birth-death-shift (BDS) processes are frequently used in stochastic modeling, with many applications in ecology and epidemiology. In particular, such processes can model evolutionary dynamics of transposable elements — important genetic markers in molecular epidemiology. Estimation of the effects of individual covariates on the birth, death, and shift rates of the process can be accomplished by analyzing patient data, but inferring these rates in a discretely and unevenly observed setting presents computational challenges. We propose a multi-type branching process approximation to BDS processes and develop a corresponding expectation maximization algorithm, where we use spectral techniques to reduce calculation of expected sufficient statistics to low dimensional integration. These techniques yield an efficient and robust optimization routine for inferring the rates of the BDS process, and apply more broadly to multi-type branching processes where rates can depend on many covariates. After rigorously testing our methodology in simulation studies, we apply our method to study intrapatient time evolution of IS*6110* transposable element, a genetic marker frequently used during estimation of epidemiological clusters of *Mycobacterium tuberculosis* infections.

1 Introduction

Continuous-time branching processes are widely used in stochastic modeling of population dynamics. Originally introduced as a mathematical model for the survival of family surnames, the tools from branching process theory have since found a breadth of applications including biology, genetics, epidemiology, quantum optics, and nuclear fission [Renshaw, 2011]. One of the most widely used classes of branching processes are birth-death (BD) processes, a simple yet flexible model for single-species population dynamics. The popularity of BD processes is in part attributable to their well-understood mathematical properties. To accurately model behavior in many applications, however, it is often necessary to consider systems with more than one species — bivariate or other multi-type processes are commonly used to model phenomena such as competition, predation, or infection [Neyman et al., 1956, Alsmeyer, 1993, Renshaw, 2011]. Multi-type branching processes form one class of models that can accommodate populations with multiple types, but these models pose considerable computational challenges for statistical inference. Our work introduces new methods to overcome these challenges, enabling likelihood-based inference in partially observed, multi-type branching processes.

Many statistically relevant quantities are available in closed form for the linear, homogeneous BD process and several of its variants, including transition probabilities, stationary distributions, and moments [Bailey, 1964, Keiding, 1975]. Further, analytical expressions of transition probabilities as series of orthogonal polynomials are known for general, nonlinear BD processes [Karlin and McGregor, 1958] and can be conveniently computed numerically [Murphy and O'donohoe, 1975, Crawford and Suchard, 2012]. The ability to compute finite-time transition probabilities enables likelihood-based inference for discretely observed or partially observed BD processes, since the observed likelihood is a function of these transition probabilities. Evaluating this likelihood is necessary in maximum likelihood estimation as well as in many Bayesian inferential procedures.

Recent work by Doss et al. [2013] and Crawford et al. [2014] introduces techniques to additionally compute conditional moments of BD sufficient statistics for linear and general birth-death-immigration processes, enabling calculation of the expected complete-data likelihood necessary in an expectation-maximization (EM) algorithm [Dempster et al., 1977].

Unfortunately, methods to evaluate finite-time transition probabilities and conditional moments are not known in the multi-type setting, and generalizing the techniques available in the single-species case is nontrivial. Solutions to the Kolmogorov equations in multi-type settings are available only for several linear, closed systems such as the immigration-death-shift process, but simple modifications such as the presence of birth events significantly complicate analysis [Puri, 1968, Renshaw, 2011]. Without these quantities, likelihood-based estimation is limited to simulation-based inference via Monte Carlo EM or MCMC [Golinelli, 2000, Golinelli et al., 2006] and asymptotic approximations, such as moment-based estimating equations [Catlin et al., 2001]. However, these approaches have shortcomings. MCMC approaches require augmenting the state space by high-dimensional latent variables and become computationally prohibitive when the state space is large. Moment-based methods are statistically less efficient than likelihood-based approaches and thus often inappropriate for smaller datasets, requiring a large number of observations to produce meaningful standard errors and confidence intervals.

In this paper, we extend the analysis of Doss et al. [2013], deriving previously unavailable numerical solutions to transition probabilities and conditional moments for discretely observed, multi-type branching processes. Modifying ideas introduced by Kendall [1948], we simplify the systems of backward equations for several relevant generating functions, and then apply the spectral approach of Lange [1982] to extract expected sufficient statistics and transition probabilities. This enables us to evaluate the observed likelihood, as well as to reduce the challenging computation of expected complete-data log-likelihood necessary in an EM algorithm to efficient evaluation of expected sufficient statistics by low-dimensional integration. Our EM algorithm can be applied in settings where the data are assumed to be generated from independent, continuous-time multi-type branching processes, observed at discrete and possibly irregularly spaced time points, whose rates can be a function of many process-specific covariates. Medical applications, for instance, commonly feature such panel data, where rates corresponding to the transmission of a disease or growth of a cell may depend on patient-specific characteristics. While similar methods have been explored for fitting continuous-time finite state-space Markov chains to panel data [Jackson, 2011, Kalbfleisch and Lawless, 1985, Lange, 1995, Lange and Minin, 2013], our method allows for multivariate and potentially infinite state-space processes.

Though our methodology applies broadly to multi-type Galton Watson branching processes, we focus attention to estimating the rates of a birth-death-shift (BDS) process. The BDS process adds the possibility of shift events to the standard BD framework, and is useful for modeling systems that allow for elements to switch locations or types — a shift is essentially a simultaneous birth and death. For example, in epidemiological applications, interaction between infected and susceptible populations can be captured as a shift event, involving a simultaneous increase and decrease in the respective populations. Spatial BDS processes have also been studied to improve Metropolis-Hastings algorithms for perfect sampling [Huber, 2012] relevant to a range of spatial statistical applications; see Illian et al. [2008] for an overview. Our motivation stems from the BDS process proposed by Rosenberg et al. [2003] to model evolution of transposons — mobile genetic elements that can replicate, die, or shift locations along the genome. Specifically, Rosenberg et al. [2003] study the within-host evolution of the IS*6110* transposon in the *Mycobacterium tuberculosis* genome via a BDS model. Accurately estimating the rates of these events is important in molecular epidemiology [Tanaka and Rosenberg, 2001].

Rosenberg et al. [2003] infer the birth, death, and shift rates of the BDS model of IS*6110*

evolution from an ongoing database of *M. tuberculosis* patients from San Francisco, but their rate estimates rely on approximate maximum likelihood estimators (MLEs) under a restrictive assumption that at most one event occurs per observation interval. This study was revisited in [Doss et al., 2013], in which the authors derive an EM algorithm for inference in discretely observed birth-death-immigration (BDI) processes amenable to high-dimensional optimization. Although this approach more realistically allows multiple events to occur per interval, the methodology in [Doss et al., 2013] is limited to the single-species setting, effectively replacing the BDS model with a simpler model that ignores particle locations and shift events. Our methodology extends the analysis of Doss et al. [2013], allowing for both the possibility of multiple events per observation interval as well as the consideration of shift events. We show that the dynamics of the BDS model can be captured in a two-type branching process framework in that transition probabilities of both processes are nearly identical. We then derive an EM algorithm for discretely observed multi-type branching processes, and rigorously assess its performance in several simulation studies. Finally, we revisit the San Francisco tuberculosis dataset, applying our algorithm to estimate rates of the IS*6110* transposon as a function of relevant covariates.

2 Methodology

2.1 Birth-death-shift model for transposable elements

Our motivation stems from a birth-death-shift process proposed by Rosenberg et al. [2003] to model evolutionary dynamics of transposable elements or transposons — genomic mobile sequence elements. Each transposon can (1) duplicate, with the new copy moving to a new genomic location; (2) shift to a different genomic position; or (3) be removed and lost from the genome, independently of all other transposons. These events occur at instantaneous rates proportional to the total transposon copy number at that time. Thus, transposons evolve according to a linear birth-death-shift (BDS) process in continuous time.

The process of transposon evolution within a host is observable by serially genotyping the organism of interest, e.g., *Mycobacterium tuberculosis* as in Rosenberg et al. [2003]. *M. tuberculosis* genome typically has between 0 and 25 copies of the IS6110 element. The number and chromosomal position of the IS*6110* element can be visualized using restriction fragment length polymorphism (RFLP). This technique entails restriction endonuclease digestion of the *M. tuberculosis* DNA which is run in an agarose gel, southern blotting and probing with a peroxidase labeled IS*6110* probe. Birth, death, and shift events are thus detectable via changes in the number and size of the bands where the IS6110 elements are located.

Estimating the rates based on observed changes at genotyping times in this experimental setup corresponds to inference in a discretely observed linear BDS process. That is, we assume each element behaves independently, and that overall rates of each event are proportional to total copy number k. Together with the time-homogeneity assumption, waiting times until occurrence of an event are distributed exponentially with rate $k\theta$, where $\theta = \lambda + \mu + \nu$. When an event occurs, the probability that it is birth, death, or shift is given by $\lambda/\theta, \nu/\theta$, and μ/θ respectively. The BDS process is therefore a continuous-time Markov chain (CTMC).

The states in our process $\widetilde{\mathbf{x}} \in \{0,1\}^S := \widetilde{\Omega}$ can be represented as binary vectors, where S is the number of possible locations transposons may occupy along the genome, 0's denote unoccupied sites, and 1's correspond to sites occupied by a transposon. Now, denote the $2^S \times 2^S$ rate matrix or infinitesimal generator corresponding to this CTMC as $\mathbf{Q} = \left\{q_{\widetilde{\mathbf{x}}_1,\widetilde{\mathbf{x}}_2}\right\}$, where $q_{\widetilde{\mathbf{x}}_1,\widetilde{\mathbf{x}}_2}$ denotes the instantaneous rate of jumping to $\widetilde{\mathbf{x}}_2$ beginning from $\widetilde{\mathbf{x}}_1$ with $\widetilde{\mathbf{x}}_1, \widetilde{\mathbf{x}}_2 \in \widetilde{\Omega}$. To write down the entries of $\mathbf{Q}$, first define $C^+(\widetilde{\mathbf{x}})$ as the set of all configurations with one additional site occupied relative

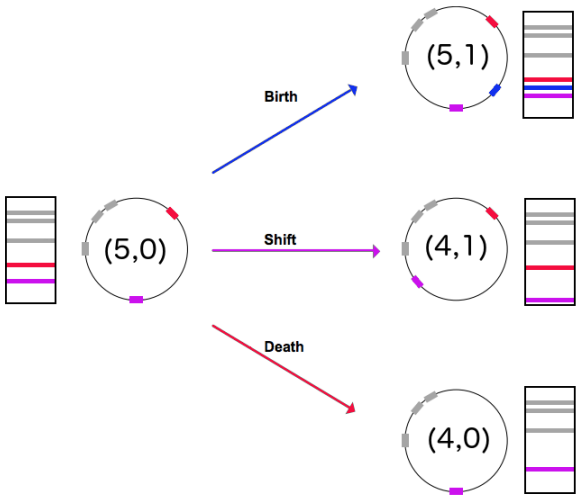

Figure 1: Illustration of the three types of transposition—birth, death, shift—along a genome, represented by circles. Transposons, depicted by filled rectangles along the circles/genomes, correspond to observable gel bands, denoted by horizontal lines in the rectangles next to each circle diagram. Numbers within each circle represent each configuration $\mathbf{X}(t)$ in the notation introduced in section 2.3. More specifically, we call the gel band on the left our initial configuration and set the number of particles of type 1 to the number of bands, 5, and the number of particles of type 2 to 0. On the right set of diagrams, a birth event keeps the number of type 1 particles intact and increments the number of type 2 particles by one, a death event changes the number of type 1 particles from five to four and keeps the number of type 2 particles at zero, and finally a shift event decreases the number of type 1 particles by one and increases the number of type 2 particles by one.

to $\widetilde{\mathbf{x}}$. Thus, $C^+(\widetilde{\mathbf{x}})$ contains states corresponding to one birth event beginning with $\widetilde{\mathbf{x}}$. Similarly $C^{\rightarrow}(\widetilde{\mathbf{x}})$ contains states where one additional site is occupied and one originally occupied site is no longer occupied, and $C^-(\widetilde{\mathbf{x}})$ contains states where one originally occupied site in $\widetilde{\mathbf{x}}$ is no longer occupied. Then $|C^+(\widetilde{\mathbf{x}}_1)| = S - k$, $|C^-(\widetilde{\mathbf{x}}_1)| = k$, and $|C^{\rightarrow}(\widetilde{\mathbf{x}}_1)| = |C^+(\widetilde{\mathbf{x}}_1)| \times |C^-(\widetilde{\mathbf{x}}_1)|$, and finally the entries of the generator $\mathbf{Q}$ are given by

$$q_{\widetilde{\mathbf{x}}_1,\widetilde{\mathbf{x}}_2} = \frac{\lambda}{|C^+(\widetilde{\mathbf{x}}_1)|}\mathbf{1}_{\{\widetilde{\mathbf{x}}_2 \in C^+(\widetilde{\mathbf{x}}_1)\}} + \frac{\nu}{|C^{\rightarrow}(\widetilde{\mathbf{x}}_1)|}\mathbf{1}_{\{\widetilde{\mathbf{x}}_2 \in C^{\rightarrow}(\widetilde{\mathbf{x}}_1)\}} + \frac{\mu}{|C^-(\widetilde{\mathbf{x}}_1)|}\mathbf{1}_{\{\widetilde{\mathbf{x}}_2 \in C^-(\widetilde{\mathbf{x}}_1)\}}. \tag{1}$$

2.2 BDS process with covariates

We are interested in inference when the data consist of m independent processes $\left\{\widetilde{\mathbf{X}}^p(t)\right\}$, $p = 1, \ldots, m$, each discretely observed at times $0 = t_{p,0} < t_{p,1} < \ldots < t_{p,n(p)}$. We assume each $\{\widetilde{\mathbf{X}}^p(t)\}$ process evolves according to a linear BDS model with per-particle instantaneous birth rate $\lambda_p \geq 0$, shift rate $\nu_p \geq 0$, and death rate $\mu_p \geq 0$. The data, observations from each process, are points in the previously defined state space, with $\widetilde{\mathbf{X}}^p(t) \in \widetilde{\Omega}$ for any fixed p and t. For example, in transposon evolution, each patient p is genotyped at $n(p)+1$ observation times, and at each given time, the 1's present in the data vector correspond to locations in the gel currently occupied by transposons. The observed data corresponding to a given process $\{\widetilde{\mathbf{X}}^p(t)\}$ can thus be collected in a $S \times [n(p)+1]$ matrix with columns corresponding to observation times, and the full observed dataset can be collected into a $S \times \sum_{p=1}^{m}[n(p)+1]$ matrix: $\mathbf{Y} = (\widetilde{\mathbf{X}}^1(t_{1,0}), \ldots, \widetilde{\mathbf{X}}^1(t_{1,n(1)}), \ldots, \widetilde{\mathbf{X}}^m(t_{m,0}), \ldots, \widetilde{\mathbf{X}}^m(t_{m,n(m)}))$.

The rates of each process are determined by a vector of c covariates $\mathbf{z}_p = (z_{p,1}, z_{p,2}, \ldots, z_{p,c}) \in \mathbb{R}^c$

through a log-linear model:

$$\log(\lambda_p) = \boldsymbol{\beta}^{\lambda} \cdot \mathbf{z_p}, \quad \log(\nu_p) = \boldsymbol{\beta}^{\nu} \cdot \mathbf{z_p}, \quad \log(\mu_p) = \boldsymbol{\beta}^{\mu} \cdot \mathbf{z_p}, \tag{2}$$

where $\boldsymbol{\beta} := (\boldsymbol{\beta}^{\lambda}, \boldsymbol{\beta}^{\nu}, \boldsymbol{\beta}^{\mu})$ are the regression coefficients and $\cdot$ represents a vector product. For instance, in an epidemiological study, these covariates may contain patient-specific disease process and demographic information.

The observed data log-likelihood is obtained by summing over transition terms of observations for each process, and summing over all processes:

$$\widetilde{\ell}_o(\mathbf{Y}; \boldsymbol{\beta}) = \sum_{p=1}^{m} \sum_{j=0}^{n(p)-1} \log \widetilde{p}_{\widetilde{\mathbf{X}}^p(t_{p,j}), \widetilde{\mathbf{X}}^p(t_{p,j+1})}(t_{p,j+1} - t_{p,j}; \lambda_p, \nu_p, \mu_p), \tag{3}$$

where $\widetilde{p}_{\widetilde{\mathbf{x}}_1, \widetilde{\mathbf{x}}_2}(t; \lambda, \nu, \mu) = \Pr_{\lambda,\nu,\mu}(\widetilde{\mathbf{X}}(t) = \widetilde{\mathbf{x}}_2 \mid \widetilde{\mathbf{X}}(0) = \widetilde{\mathbf{x}}_1)$ denotes a transition probability of the BDS process. We are interested in computing the maximum likelihood estimates (MLEs) of parameters $\boldsymbol{\beta}$ of the BDS process. Notice that if the transition probabilities were available for given λ, ν, μ, and t values, one could maximize the likelihood in (3) using standard off-the-shelf optimization procedures. However, due to the large state space of all possible configurations of occupied sites, analysis of these transition probabilities is intractable. To approximate the BDS model likelihood above, we introduce a two-type branching process such that computationally tractable transition probabilities of this process are numerically close to the transition probabilities of the BDS model over any observation interval. The following sections detail the correspondence between the BDS model and the two-type branching process, and develops methodology for inference in the branching process framework.

2.3 Reducing the state space

The size of the original state space $|\widetilde{\Omega}| = 2^S$ quickly becomes unmanageable as S grows so that analysis using the rate matrix defined in (1) becomes unwieldy for all but small values of S. Previous work by Doss et al. [2013] addresses this issue by collapsing the state space to one dimension, distilling the data to copy number counts at each observation time. In this simplified setting, they develop tools for inference in a discretely observed birth-death-immigration framework. However, this approximate model ignores particle shifts which do not affect the total copy number, rendering the shift rate unidentifiable. Further, collapsing the state space in this way violates the Markov assumption in the BDS model. In particular, waiting times between birth and death events are exponentially distributed under the model in Doss et al. [2013], but under the BDS model with shift events, the waiting time between a birth and death no longer follows an exponential distribution.

Instead of ignoring shifts, we propose a reduction of the state space into a two-dimensional representation $\Omega \in \mathbb{N} \times \mathbb{N}$. Elements of this reduced space are pairs $\mathbf{X}(t) = (x_{old}, x_{new}) \in \Omega$ tracking the number of originally occupied and newly occupied sites at the end of each observation interval. As an example, assume six particles are present initially at time t_0, and a shift and a birth occur before the first observation t_1, and a death occurs before a second observation at t_2. When considering the first observation interval $[t_0, t_1)$, we have $\{\mathbf{X}(t_0) = (6, 0), \mathbf{X}(t_1) = (5, 2)\}$. When computing the next transition probability over $[t_1, t_2)$, we now have $\{\mathbf{X}(t_1) = (7, 0), \mathbf{X}(t_2) = (6, 0)\}$, since all seven of the particles at t_1, now the left endpoint of the observation interval, now become the initial population. This seemingly inconsistent definition of the state at $\mathbf{X}(t_1)$ is not a problem: we will see that all necessary computations occur separately on disjoint intervals, so that our reduced representation of the original process needs only to be defined consistently for any given pair of consecutive observations.

Formally, this state space transformation is a mapping $\psi : \widetilde{\Omega} \times \widetilde{\Omega} \to \Omega \times \Omega$ on consecutive pairs of observations in $\widetilde{\Omega}$ to the reduced state space that can be computed $\psi : \left\{\widetilde{\mathbf{X}}(t_1), \widetilde{\mathbf{X}}(t_2)\right\} \mapsto \{(a,0),(b,c)\} = \{\mathbf{X}(t_1), \mathbf{X}(t_2)\}$, where $a = \sum_{j=1}^{S} \widetilde{X}_j(t_1)$ is the total number of initially occupied sites in $\widetilde{\mathbf{X}}(t_1)$, $b = \sum_{j=1}^{S} \mathbf{1}_{\{\widetilde{X}_j(t_2)=\widetilde{X}_j(t_1)\}} \widetilde{X}_j(t_1)$ is the number of initially occupied sites that remain occupied, and $c = \sum_{j=1}^{S} \mathbf{1}_{\{\widetilde{X}_j(t_2)-\widetilde{X}_j(t_1)=1\}}$ is the number of newly occupied sites in $\widetilde{\mathbf{X}}(t_2)$ not present in $\widetilde{\mathbf{X}}(t_1)$.

Note that while ψ significantly reduces the size of the state space, the mapping discards information about specific particle locations, which is uninformative to inferring birth, death, and shift rates due to symmetry induced by particle independence. The number of changes in locations between observations — the data relevant to our estimation task — is preserved in the image of ψ.

2.4 Modeling via two-type branching process

Working now in the space Ω, we can treat x_{old} and x_{new} as particle types in a two-type branching process. Let $a_j(k,l)$ be the rate of producing k type 1 particles and l type 2 particles, beginning with one type j particle, $j = 1, 2$. Then the nonzero rates defining the two-type branching process corresponding to the birth-death-shift model are given by

$$\begin{aligned} &a_1(1,1) = \lambda, \qquad a_1(0,1) = \nu, \qquad a_1(0,0) = \mu, \qquad a_1(1,0) = -(\lambda+\nu+\mu), \\ &a_2(0,2) = \lambda, \qquad a_2(0,1) = -(\lambda+\mu), \qquad a_2(0,0) = \mu. \end{aligned} \tag{4}$$

This characterization enables us to apply a generating function approach to calculate transition probabilities of the process. Defining $X_j(t)$, the number of particles of type j at time t, we consider the generating function

$$\phi_{jk}(t, s_1, s_2) = \mathrm{E}\left(s_1^{X_1(t)} s_2^{X_2(t)} \mid X_1(0) = j, X_2(0) = k\right) = \sum_{l=0}^{\infty}\sum_{m=0}^{\infty} p_{(j,k),(l,m)}(t) s_1^l s_2^m. \tag{5}$$

Using the Kolmogorov backward equations, we derive equations and a closed form solution for ϕ_{jk} (see Appendix A). Although an analytical expression is available, it involves special functions that are in practice are often unstable. Instead, we simplify the backward equations so that evaluating ϕ_{jk} only requires solving a single linear ordinary differential equation, which is easily accomplished using standard Runge-Kutta methods [Butcher, 1987].

With ϕ_{jk} available, we see from (5) that the transition probabilities $p_{(j,k),(l,m)}(t)$ can then be obtained by differentiating and setting $s_1, s_2 = 0$, but without an analytical expression for these derivatives, repeated numerical differentiation is inefficient and numerically unstable. Instead, we follow Lange [1982] and Doss et al. [2013] and map our domain $[0,1] \times [0,1]$ to the boundary of a unit circle in the complex plane by setting $s_1 = e^{2\pi i w_1}, s_2 = e^{2\pi i w_2}$. Under this change of variables, the generating function becomes a Fourier series

$$\phi_{jk}\left(t, e^{2\pi i w_1}, e^{2\pi i w_2}\right) = \sum_{l,m=0}^{\infty} p_{(j,k),(l,m)}(t) e^{2\pi i l w_1} e^{2\pi i m w_2}.$$

Applying a Riemann sum approximation to the integral corresponding to coefficients given by the

Fourier inversion formula, we can compute the transition probabilities using

$$\begin{aligned} p_{(j,k),(l,m)}(t) &= \int_0^1 \int_0^1 \phi_{jk}(t, e^{2\pi i w_1}, e^{2\pi i w_2}) e^{-2\pi i l w_1} e^{-2\pi i m w_2} dw_1 dw_2 \\ &\approx \frac{1}{N^2} \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} \phi_{jk}(t, e^{2\pi i u/N}, e^{2\pi i v/N}) e^{-2\pi i l u/N} e^{-2\pi i m v/N}. \end{aligned} \tag{6}$$

Choice of a larger N leads to a finer and thus more accurate Riemann sum approximation of the integral, and also allows us to compute transition probabilities to and from a larger total particle population of either type. The Fast Fourier transform (FFT) enables efficient computation of these coefficients [Henrici, 1979], and in our application and simulation studies, we find that a grid size as small as $N = 16$ yields accurate results. With transition probabilities available, we may closely approximate $\widetilde{\ell}_o(\mathbf{Y}; \boldsymbol{\beta})$ by the branching process likelihood

$$\ell_o(\mathbf{Y}; \boldsymbol{\beta}) = \sum_{p=1}^{m} \sum_{j=0}^{n(p)-1} \log p_{\mathbf{X}^p(t_{p,j}), \mathbf{X}^p(t_{p,j+1})}(t_{p,j+1} - t_{p,j}; \lambda_p, \nu_p, \mu_p), \tag{7}$$

so that maximizing the observed likelihood in (3) gives approximately the same parameter estimates as maximizing (7).

2.5 EM algorithm for the BDS process

With transition probabilities of the process available, it is already possible to produce MLEs of the covariate effects associated with birth, death, and shift rates by numerical maximization of the observed likelihood. However, an EM algorithm approach often outperforms off-the-shelf optimization procedures in missing data problems, offering a significantly faster and more robust solution. Let $\ell_c(\mathbf{X}, \boldsymbol{\beta})$ denote the complete data log-likelihood, $\mathbf{X}$ the complete data, and $\mathbf{Y}$ the available observations. The EM algorithm begins with an initial parameter estimate $\boldsymbol{\beta}_0$, and then at each j^{th} iteration, updates the estimate by setting

$$\boldsymbol{\beta}_j = \underset{\boldsymbol{\beta}}{\operatorname{argmax}} \, \mathrm{E}_{\boldsymbol{\beta}_{j-1}} \left[\ell_c(\mathbf{X}, \boldsymbol{\beta}) \mid \mathbf{Y} \right]. \tag{8}$$

Each iteration involves a computation of the expectation term called the *E-step*, followed by a maximization of the expectation called the *M-step*.

2.5.1 E-step

The fully observed BDS process is a continuous-time Markov chain, so its complete-data log-likelihood can be written as

$$\ell_c(\mathbf{X}; \boldsymbol{\beta}) = \sum_{p=1}^{m} \left[b_p \log \lambda_p + f_p \log \nu_p + d_p \log \mu_p - (\lambda_p + \mu_p + \nu_p) \sum_{k=0}^{\infty} k \tau_p(k) + \sum_{k=0}^{\infty} \log \tau_p(k) \right], \tag{9}$$

where $\tau_p(k)$ is the total time process $\mathbf{X}^p(t)$ spends with total copy number $X_1^p(t) + X_2^p(t) = k$, b_p is the total number of births, f_p the number of shifts, and d_p the number of deaths for each patient $p = 1, \ldots, m$ — these quantities are the complete data sufficient statistics [Guttorp, 1995]. Notice the final term in (9) is constant with respect to the parameters. We see that in order to obtain the expected complete-data log-likelihood, we need to calculate only expected births — $\mathrm{E}_{\boldsymbol{\beta}} \left[b_p \mid \mathbf{Y} \right]$,

shifts — $\mathrm{E}_{\boldsymbol{\beta}}[f_p \mid \mathbf{Y}]$, deaths — $\mathrm{E}_{\boldsymbol{\beta}}[d_p \mid \mathbf{Y}]$, and particle time — $\mathrm{E}_{\boldsymbol{\beta}}[R_p \mid \mathbf{Y}]$, where the last quantity is defined as

$$\mathrm{E}_{\boldsymbol{\beta}}[R_p \mid \mathbf{Y}] := \mathrm{E}_{\boldsymbol{\beta}}\left[\int_{t_{p,0}}^{t_{p,n(p)}} X_1(s) + X_2(s)ds \mid \mathbf{Y}\right] = \mathrm{E}_{\boldsymbol{\beta}}\left[\sum_{k=0}^{\infty} k\tau_p(k) \mid \mathbf{Y}\right].$$

By independence of the p processes and linearity of expectations, each expectation breaks into sums of expectations over the observation intervals. Further, by homogeneity, it suffices to be able to calculate the quantities

$$\begin{aligned}
e^{+}_{jk,lm}(t) &= \mathrm{E}\left[b_{p,t} \mid \mathbf{X}^p(0) = (j,k), \mathbf{X}^p(t) = (l,m)\right], \\
e^{\rightarrow}_{jk,lm}(t) &= \mathrm{E}\left[f_{p,t} \mid \mathbf{X}^p(0) = (j,k), \mathbf{X}^p(t) = (l,m)\right], \\
e^{-}_{jk,lm}(t) &= \mathrm{E}\left[d_{p,t} \mid \mathbf{X}^p(0) = (j,k), \mathbf{X}^p(t) = (l,m)\right], \\
e^{*}_{jk,lm}(t) &= \mathrm{E}\left[R_{p,t} \mid \mathbf{X}^p(0) = (j,k), \mathbf{X}^p(t) = (l,m)\right],
\end{aligned}$$

for all non-negative integers j, k, l, m. Dependence of these quantities on rates λ_p, ν_p, μ_p is suppressed in the notation here for simplicity. As noticed by Minin and Suchard [2008] and Doss et al. [2013], it is easier to work via the restricted moments

$$\begin{aligned}
m^{+}_{jk,lm}(t) &= \mathrm{E}\left[b_{p,t}1_{\{\mathbf{X}^p(t)=lm\}} \mid \mathbf{X}^p(0) = (j,k)\right] = \sum_{n=0}^{\infty} n q^{+}_{jk,lm}(n,t), \\
m^{\rightarrow}_{jk,lm}(t) &= \mathrm{E}\left[f_{p,t}1_{\{\mathbf{X}^p(t)=lm\}} \mid \mathbf{X}^p(0) = (j,k)\right] = \sum_{n=0}^{\infty} n q^{\rightarrow}_{jk,lm}(n,t), \\
m^{-}_{jk,lm}(t) &= \mathrm{E}\left[d_{p,t}1_{\{\mathbf{X}^p(t)=lm\}} \mid \mathbf{X}^p(0) = (j,k)\right] = \sum_{n=0}^{\infty} n q^{-}_{jk,lm}(n,t), \\
m^{*}_{jk,lm}(t) &= \mathrm{E}\left[R_{p,t}1_{\{\mathbf{X}^p(t)=lm\}} \mid \mathbf{X}^p(0) = (j,k)\right] = \int_{x=0}^{\infty} x \mathrm{d} q^{*}_{jk,lm}(x,t),
\end{aligned}$$

where

$$\begin{aligned}
q^{*}_{jk,lm}(x,t) &= \Pr[R_{p,t} \le x, \mathbf{X}^p(t) = (l,m) \mid \mathbf{X}^p(0) = (j,k)], \\
q^{+}_{jk,lm}(n,t) &= \Pr[b_{p,t} = n, \mathbf{X}^p(t) = (l,m) \mid \mathbf{X}^p(0) = (j,k)],
\end{aligned}$$

and $q^{\rightarrow}, q^{-}$ are defined analogously. The conditional expectations can then be recovered after dividing by transition probabilities, i.e.

$$e^{+}_{jk,lm}(t) = m^{+}_{jk,lm}(t)/p_{jk,lm}(t).$$

These restricted moments can be computed with a similar approach used to obtain transition probabilities. We begin by defining the pseudo-generating functions: for expected births, let

$$g^{+}_{jk,lm}(r,t) = \sum_{n=0}^{\infty} q^{+}_{jk,lm}(n,t) r^n.$$

Ignoring notational dependence on individual patients for simplicity, we define the joint generating function

$$\begin{aligned}
&H^{+}_{jk}(r, s_1, s_2, t) = \mathrm{E}\left[r^{b_t} s_1^{X_1(t)} s_2^{X_2(t)} \mid \mathbf{X}(0) = (j,k)\right] \\
&= \sum_l \sum_m \sum_n \Pr\left[b_t = n, \mathbf{X}(t) = (k,l) \mid \mathbf{X}(0) = (j,k)\right] r^n s_1^l s_2^m = \sum_l \sum_m g^{+}_{jk,lm}(r,t) s_1^l s_2^m.
\end{aligned}$$

Pseudo-generating functions for shifts and deaths are defined analogously, and the pseudo-generating function for particle time is defined as

$$H^*_{jk}(r, s_1, s_2, t) = \sum_l \sum_m \int_{x=0}^{\infty} e^{-rx} \mathrm{d}q^*_{jk,lm}(x, t) s_1^l s_2^m := \sum_l \sum_m V_{jk,lm}(r, t) s_1^l s_2^m,$$

where $V_{jk,lm}(r; t) = \int_0^\infty e^{-rx} \mathrm{d}q^*_{jk,lm}(x; t)$ is the Laplace-Stieltjes transform of $q^*_{jk,lm}(x; t)$. In each case we can define series whose coefficients are our quantities of interest by partial differentiation:

$$G^+_{jk}(s_1, s_2, t) = \frac{d}{dr} H^+_{jk}(r, s_1, s_2, t)\bigg|_{r=1} = \sum_l \sum_m \left[\sum_n n q^+_{jk,lm}(n, t) \right] s_1^l s_2^m = \sum_l \sum_m m^+_{jk,lm}(t) s_1^l s_2^m. \tag{10}$$

$G^{\rightarrow}_{jk}$ and G^-_{jk} are defined analogously, and the expression for particle time is instead differentiated at $r = 0$:

$$G^*_{jk}(s_1, s_2, t) = \frac{d}{dr} H^*_{jk}(r, s_1, s_2, t)\bigg|_{r=0} = \sum_l \sum_m \left[\int_{x=0}^{\infty} x \mathrm{d}q^*_{jk,lm}(x, t) \right] s_1^l s_2^m = \sum_l \sum_m m^*_{jk,lm}(t) s_1^l s_2^m. \tag{11}$$

We see that given expressions for H^+_{jk}, $H^{\rightarrow}_{jk}$, H^-_{jk}, and H^*_{jk}, the coefficients corresponding to moments $m^+_{jk,lm}$, $m^{\rightarrow}_{jk,lm}$, $m^-_{jk,lm}$, $m^*_{jk,lm}$ can then be numerically computed using FFT analogously to (6) by replacing ϕ_{jk} with the corresponding G_{jk} functions. For notational simplicity, we use G_{jk} when referring collectively to G^+_{jk}, $G^{\rightarrow}_{jk}$, G^-_{jk}, and G^*_{jk}, and similarly define H_{jk}.

Having reduced our task to computing H_{jk}, we define $H_1 := H_{10}(r, s_1, s_2, t)$ and $H_2 := H_{01}(r, s_1, s_2, t)$. By independence of particles in the branching process, we have $H_{jk} = H_1^j H_2^k$. In all four cases, H_2 is analytically available, and we derive an ordinary differential equation for H_1, summarized in the theorem below. We present the result for a branching process with rates corresponding to the birth-death-shift model, but such systems of equations are available for an arbitrary time-homogeneous multi-type branching process.

Theorem 2.1 *Let $\{X_t\}$ be a two-type branching defined by the rates in* (A-4)*. Denote particle time and the number of births, shifts, and deaths over the interval $[0, t)$ by R_t, b_t, f_p, and d_t respectively. Define the generating functions corresponding to births as*

$$H^+_1(r, s_1, s_2, t) = E\left[r^{b_t} s_1^{X_1(t)} s_2^{X_2(t)} \mid \mathbf{X}(0) = (1, 0) \right] \text{ and}$$
$$H^+_2(r, s_1, s_2, t) = E\left[r^{b_t} s_1^{X_1(t)} s_2^{X_2(t)} \mid \mathbf{X}(0) = (0, 1) \right].$$

Then

$$H^+_2 = y_b + \left[\frac{-\lambda r}{2\lambda r y_b - \lambda - \mu} + \left(\frac{1}{s_2 - y_b} + \frac{\lambda r}{2\lambda r y_b - \lambda - \mu} \right) e^{-(2y_b \lambda r - \lambda - \mu)t} \right]^{-1},$$

where $y_b = (\lambda + \mu + \sqrt{\lambda^2 + 2\lambda\mu + \mu^2 - 4\lambda\mu r})/(2\lambda r)$, and H^+_1 satisfies the following differential equation:

$$\frac{d}{dt} H^+_1(t, s_1, s_2, r) = \lambda r H^+_1 H^+_2 + \nu H^+_2 + \mu - (\lambda + \mu + \nu) H^+_1, \tag{12}$$

subject to initial condition $H_1(r, s_1, s_2, 0) = s_1$.

The analogous generating functions for shifts, deaths, and particle time satisfy the following equations:

$$H_2^-(t,s_1,s_2,r) = y_d + \left[\frac{-\lambda}{2\lambda y_d - \lambda - \mu} + \left(\frac{1}{s_2 - y_d} + \frac{\lambda}{2\lambda y_d - \lambda - \mu}\right) e^{-(2y_d\lambda - \lambda - \mu)t}\right]^{-1},$$
$$H_2^{\rightarrow}(t,s_1,s_2,r) = 1 + \left[\frac{\lambda}{\mu - \lambda} + (\frac{1}{s_2 - 1} + \frac{\lambda}{\lambda - \mu})e^{(\mu-\lambda)t}\right]^{-1},$$
$$H_2^*(t,s_1,s_2,r) = y_* + \left[\frac{-\lambda}{2\lambda y_* - \lambda - \mu - r} + \left(\frac{1}{s_2 - y_*} + \frac{\lambda}{2\lambda y_* - \lambda - \mu - r}\right) e^{-(2y_*\lambda - \lambda - \mu - r)t}\right]^{-1},$$

$$\frac{d}{dt}H_1^-(t,s_1,s_2,r) = \lambda H_1^- H_2^- + \nu H_2^- + \mu r - (\lambda + \mu + \nu)H_1^-,$$
$$\frac{d}{dt}H_1^{\rightarrow}(t,s_1,s_2,r) = \lambda H_1^{\rightarrow} H_2^{\rightarrow} + \nu r H_2^{\rightarrow} + \mu - (\lambda + \mu + \nu)H_1^{\rightarrow},$$
$$\frac{d}{dt}H_1^*(t,s_1,s_2,r) = \lambda H_1^* H_2^* + \nu H_2^* + \mu - (\lambda + \mu + \nu + r)H_1^*,$$

where $y_d = (\lambda + \mu + \sqrt{\lambda^2 + 2\lambda\mu + \mu^2 - 4\lambda\mu r})/(2\lambda)$, $y_* = (\lambda + \mu + r + \sqrt{(\lambda + \mu + r)^2 - 4\lambda\mu})/(2\lambda)$, *and* $H_1^-(r,s_1,s_2,0) = H_1^{\rightarrow}(r,s_1,s_2,0) = H_1^*(r,s_1,s_2,0) = s_1$.

Proof The derivations are included in the Appendix B, using the birth equations (B-1) as a detailed example. The other systems follow analogous derivations.

This theorem shows that for each of the necessary sufficient statistics, computations for H_{jk} are essentially reduced to solving a single ordinary differential equation. As discussed for transition probabilities, this is easily accomplished using Runge-Kutta methods, which in practice offer more numerical stability than working with solutions obtained analytically by integrating the ODE. Because we can evaluate H_{jk}, we can also easily differentiate H_{jk} numerically, yielding access to numerical solutions to G_{jk}.

To summarize, with $H_{jk}^+, H_{jk}^-, H_{jk}^{\rightarrow}, H_{jk}^*$ now available, we may obtain the restricted moments by computing the coefficients in the power series $G_{jk}^+, G_{jk}^-, G_{jk}^{\rightarrow}, G_{jk}^*$. These coefficients are recovered using a Riemann approximation to the Fourier inversion formula analogous to formula (6). For instance,

$$\begin{aligned} m_{(jk),(lm)}^+(t) &= \int_0^1 \int_0^1 G_{jk}^+\left(t, e^{2\pi i t_1}, e^{2\pi i t_2}\right) e^{-2\pi i l t_1} e^{-2\pi i m t_2} dt_1 dt_2 \\ &\approx \frac{1}{N^2} \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} G_{jk}^+\left(t, e^{2\pi i u/N}, e^{2\pi i v/N}\right) e^{-2\pi i l u/N} e^{-2\pi i m v/N}. \end{aligned}$$

We are thus able to compute all necessary quantities appearing in the expected complete-data log-likelihood $\mathrm{E}_{\widetilde{\boldsymbol{\beta}}}\left[\ell_c(\mathbf{X}, \boldsymbol{\beta}) \mid \mathbf{Y}\right]$. Recall that sufficient statistics for each patient b_p, f_p, d_p, and R_p break up over intervals: i.e. the total number of births b_p is equal to the sum of the number of births over each disjoint interval $[t_{p,j-1}, t_{p,j})$, with $j = 1, \ldots, n(p)$. Further, by the Markov property, the conditional expectation of the number births over an interval $[t_1, t_2)$ given $\mathbf{Y}$ depends only on the states of the process at the endpoints of the interval:

$$\mathrm{E}\left[b_{p,t_2-t_1} \mid \mathbf{Y}\right] = \mathrm{E}\left[b_{p,t_2-t_1} | \mathbf{X}^p(t_1), \mathbf{X}^p(t_2)\right] = e_{\mathbf{X}^p(t_1),\mathbf{X}^p(t_2)}^+(t_2 - t_1) = \frac{m_{\mathbf{X}^p(t_1),\mathbf{X}^p(t_2)}^+(t_2 - t_1)}{p_{\mathbf{X}^p(t_1),\mathbf{X}^p(t_2)}(t_2 - t_1)}, \quad (13)$$

and the same is true for the other sufficient statistics. Therefore, for each process p,

$$\begin{aligned}
\mathrm{E}_{\widetilde{\boldsymbol{\beta}}}\left[b_p \mid \mathbf{Y}\right] &= \sum_{i=1}^{n(p)} e^{+}_{\mathbf{X}^p(t_{p,i-1}),\mathbf{X}^p(t_{p,i})}(t_{p,i-1}-t_{p,i};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p),\\
\mathrm{E}_{\widetilde{\boldsymbol{\beta}}}\left[f_p \mid \mathbf{Y}\right] &= \sum_{i=1}^{n(p)} e^{\rightarrow}_{\mathbf{X}^p(t_{p,i-1}),\mathbf{X}^p(t_{p,i})}(t_{p,i-1}-t_{p,i};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p), \text{ and}\\
\mathrm{E}_{\widetilde{\boldsymbol{\beta}}}\left[d_p \mid \mathbf{Y}\right] &= \sum_{i=1}^{n(p)} e^{-}_{\mathbf{X}^p(t_{p,i-1}),\mathbf{X}^p(t_{p,i})}(t_{p,i-1}-t_{p,i};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p),
\end{aligned} \tag{14}$$

with $\log\left(\widetilde{\lambda}_p\right) = \widetilde{\boldsymbol{\beta}}^{\lambda}\cdot\mathbf{z_p}, \log\left(\widetilde{\nu}_p\right) = \widetilde{\boldsymbol{\beta}}^{\nu}\cdot\mathbf{z_p}, \log\left(\widetilde{\mu}_p\right) = \widetilde{\boldsymbol{\beta}}^{\mu}\cdot\mathbf{z_p}$ similarly to equation (2). Finally, combining (14), (13), and (9), the expected complete-data log likelihood up to a constant is equal to

$$\begin{aligned}
\mathrm{E}_{\widetilde{\boldsymbol{\beta}}}\left[\ell_c(\mathbf{X},\boldsymbol{\beta}) \mid \mathbf{Y}\right] \propto \sum_{p=1}^{m}\Bigg\{\sum_{j=1}^{n(p)}\Bigg[&\frac{m^{+}_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}{p_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}\log\lambda_p\\
&+\frac{m^{\rightarrow}_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}{p_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}\log\nu_p\\
&+\frac{m^{-}_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}{p_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}\log\mu_p\\
&-\frac{m^{*}_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}{p_{\mathbf{X}^p(t_{p,j-1}),\mathbf{X}^p(t_{p,j})}(t_{p,j}-t_{p,j-1};\widetilde{\lambda}_p,\widetilde{\nu}_p,\widetilde{\mu}_p)}(\lambda_p+\mu_p+\nu_p)\Bigg]\Bigg\}.
\end{aligned} \tag{15}$$

2.5.2 M-step

To complete an M-step, we use an efficient Newton-Raphson algorithm to maximize the expectation $g(\boldsymbol{\beta}) = E_{\widetilde{\boldsymbol{\beta}}}\left[\ell_c(\mathbf{X},\boldsymbol{\beta}) \mid \mathbf{Y}\right]$. Each Newton-Raphson step recursively updates parameters using the following equation:

$$\boldsymbol{\beta}_{new} = \boldsymbol{\beta}_{cur} - \left[\mathbf{H}g(\boldsymbol{\beta}_{cur})\right]^{-1}\nabla g(\boldsymbol{\beta}_{cur}), \tag{16}$$

where ∇g denotes the gradient vector and $\mathbf{H}g$ denotes the Hessian matrix of $g(\boldsymbol{\beta})$. Fortunately, compact analytical forms for these quantities are available. First, we collect complete data sufficient statistics across processes into the following vectors:

$$\begin{aligned}
\mathbf{U}^T &= \left(E_{\widetilde{\boldsymbol{\beta}}}\left[b_{1,t_{1,n(1)}}|\mathbf{Y}\right],\ldots,E_{\widetilde{\boldsymbol{\beta}}}\left[b_{m,t_{m,n(m)}}|\mathbf{Y}\right]\right), \quad \mathbf{V}^T = \left(E_{\widetilde{\boldsymbol{\beta}}}\left[f_{1,t_{1,n(1)}}|\mathbf{Y}\right],\ldots,E_{\widetilde{\boldsymbol{\beta}}}\left[f_{m,t_{m,n(m)}}|\mathbf{Y}\right]\right),\\
\mathbf{D}^T &= \left(E_{\widetilde{\boldsymbol{\beta}}}\left[d_{1,t_{1,n(1)}}|\mathbf{Y}\right],\ldots,E_{\widetilde{\boldsymbol{\beta}}}\left[d_{m,t_{m,n(m)}}|\mathbf{Y}\right]\right), \quad \mathbf{P}^T = \left(E_{\widetilde{\boldsymbol{\beta}}}\left[R_{1,t_{1,n(1)}}|\mathbf{Y}\right],\ldots,E_{\widetilde{\boldsymbol{\beta}}}\left[R_{m,t_{m,n(m)}}|\mathbf{Y}\right]\right).
\end{aligned}$$

If we aggregate covariate vectors for each process in a $c\times p$ matrix $\mathbf{Z} = (\mathbf{z_1},\ldots,\mathbf{z_m})$ and process-specific rates into vectors $\boldsymbol{\lambda} = (\lambda_1,\ldots,\lambda_m), \boldsymbol{\nu} = (\nu_1,\ldots,\nu_m), \boldsymbol{\mu} = (\mu_1,\ldots,\mu_m)$, then the gradient and Hessian can be expressed as

$$\nabla g(\boldsymbol{\beta}) = \left(-\mathbf{Z}^T\left[\mathrm{diag}(\mathbf{P})\boldsymbol{\lambda}+\mathbf{U}\right], -\mathbf{Z}^T\left[\mathrm{diag}(\mathbf{P})\boldsymbol{\nu}+\mathbf{V}\right], -\mathbf{Z}^T\left[\mathrm{diag}(\mathbf{P})\boldsymbol{\mu}+\mathbf{D}\right]\right), \tag{17}$$

$$\mathbf{Hg}(\boldsymbol{\beta}) = \begin{pmatrix} -\mathbf{Z}^T\mathrm{diag}(\mathbf{P})\mathrm{diag}(\boldsymbol{\lambda})\mathbf{Z} & \mathbf{0} & \mathbf{0}\\ \mathbf{0} & -\mathbf{Z}^T\mathrm{diag}(\mathbf{P})\mathrm{diag}(\boldsymbol{\nu})\mathbf{Z} & \mathbf{0}\\ \mathbf{0} & \mathbf{0} & -\mathbf{Z}^T\mathrm{diag}(\mathbf{P})\mathrm{diag}(\boldsymbol{\mu})\mathbf{Z}\end{pmatrix}. \tag{18}$$

The derivation of these expressions is parallel to those presented in [Doss et al., 2013]. In our experience the M-step generally converges in fewer than ten Newton-Raphson steps. Availability of closed form solutions (17) and (18) yields very fast execution of each Newton-Raphson step, making the computational cost of the M-step negligible compared to the E-step. Note that computing the M-step using only one Newton-Raphson step rather than iterating until convergence is sufficient to guarantee the ascent property of the EM algorithm [Lange, 1995], but we execute multiple steps because this strategy does not slow down our algorithm.

2.5.3 Accelerating E-step calculations for intervals with no change

In our birth-death-shift application, we may avoid the relatively costly E-step calculations for some intervals by approximating the probability of observing no changes with the probability that no event occurs in the underlying complete process. This approximation is not necessary in our algorithm, but can lead to gains in computational efficiency in settings such as our application where many intervals feature no observed changes.

It is very unlikely that events occur in a time interval $[t_1, t_2)$ yet no change is observed so that $\mathbf{X}(t_1) = \mathbf{X}(t_2)$. For instance, if 12 elements are present initially and a death followed by a birth occur, then we almost always observe $\mathbf{X}(t_1) = (12, 0), \mathbf{X}(t_2) = (11, 1)$ unless the element added by the birth occupies the *exact* location that was previously occupied by the element that dies. This scenario would leave the observed state unchanged, $\mathbf{X}(t_1) = \mathbf{X}(t_2) = (12, 0)$, but has exceedingly low probability: the already small but non-negligible probability that more than one event occurs is then multiplied by $1/(S - 11)$, the probability of the birth occurring in a specific location (recall S is very large). Therefore, it is numerically accurate to treat intervals with no observed changes as if no changes in the latent continuous-time process occur. In this case, the transition probability is easily calculated, given by the tail of an exponential distribution

$$p_{(12,0),(12,0)}(t_2 - t_1) = e^{-12(\lambda+\mu+\nu)(t_2-t_1)}. \tag{19}$$

In addition to efficient closed-form transition probability calculation, the expected sufficient statistics necessary for the E-step are known in this setting. If no events occur, we know that $e^{+}_{(k,0),(k,0)}(t) = e^{\rightarrow}_{(k,0),(k,0)}(t) = e^{-}_{(k,0),(k,0)}(t) = 0$, and that the expected particle time is $e^{*}_{(k,0),(k,0)}(t) = kt$. This is not only faster computationally but also more numerically stable, avoiding the division of numerically calculated restricted moments by numerically calculated transition probabilities. We verify this efficient implementation in our simulation studies, as illustrated in Figure 4.

Equation (19) is the same formula that Rosenberg et al. [2003] used to compute the probability of no event under their frequent monitoring (FM) method, but it is important to note that this approximation is *not* accurate for the other probabilities under FM; see Figure 2. Continuing our previous example, consider the probability of transitioning from $(12, 0)$ to $(11, 1)$. Under the FM approximation, such a transition can only happen through a shift event with probability $p_{(12,0),(11,1)}(t_2 - t_1) = (\nu)/(\lambda + \nu + \mu)e^{-k(\lambda+\mu+\nu)(t_2-t_1)}$. However, as we have discussed above, a birth followed by death also leads to observing $\mathbf{X}(t_2) = (11, 1)$ in almost all cases. FM assigns zero probability density to such event histories, while our method does not ignore non-negligible contribution of these trajectories to $p_{(12,0),(11,1)}(t_2 - t_1)$. Additionally, probabilities $p_{(j,k),(l,m)}(t_2 - t_1)$ are set to zero under FM when for all j, k, l, m where $|j - k| > 1$ or $|l - m| > 1$, preventing the use of all available data during inference. This is later illustrated in Figure 3.

2.6 Implementation

We implement our algorithm in the form of R package `bdsem`, available at `https://github.com/jasonxu90/bdsem`. The EM algorithm implementation relies on numerical solutions to differential

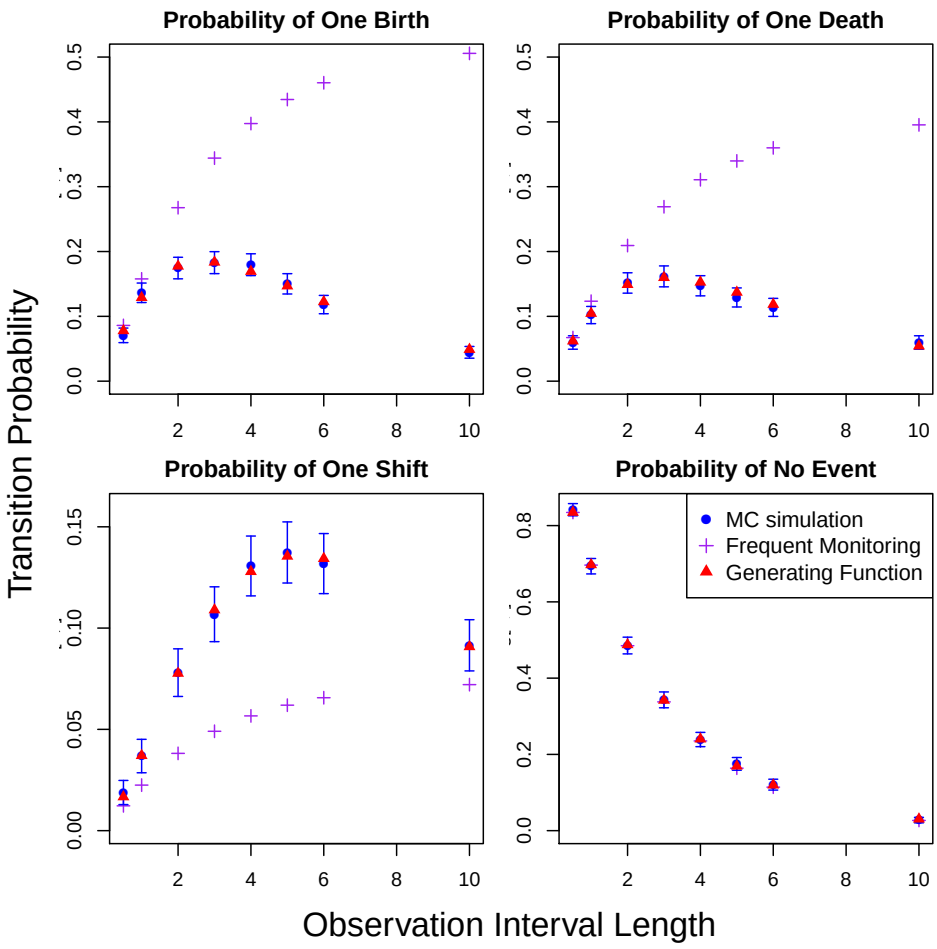

Figure 2: Transition probability approximations. BDS transition probabilities are approximated with two methods — the FM method, shown with magenta crosses, and the generating function method, shown with red triangles. We depict Monte Carlo estimates of the BDS transition probabilities with blue circles; vertical blue segments indicate their corresponding Monte Carlo confidence intervals.

equations in package `deSolve`, and accommodates panel data settings with unevenly spaced discrete observations. Our package also includes functions for MLE inference using other methods, as well as code for simulating from the BDS process. The software is accompanied by a vignette that steps through simplified versions of all simulation studies included in this paper.

2.7 Comparison with frequent monitoring

We begin with several simulation experiments assessing the validity of our algorithms. The first simulation study checks whether transition probabilities calculated using our generating function method for the two-type branching process as described in (6) coincide with those of the BDS model. We compare these computations to Monte Carlo estimates of these probabilities obtained from simulated trajectories from the birth-death-shift model, and also include a comparison to the FM method presented in [Rosenberg et al., 2003].

The FM model allows at most one event to occur per interval. Thus, over an observation interval

$[t_i, t_{i+1})$ beginning with k particles, the probabilities of a birth, death, and shift have closed forms $(\lambda/\theta)e^{-k\theta(t_{i+1}-t_i)}$, $(\mu/\theta)e^{-k\theta(t_{i+1}-t_i)}$, and $(\nu/\theta)e^{-k\theta(t_{i+1}-t_i)}$ respectively, where $\theta = \lambda + \nu + \mu$. The probability of no event occurring is given by $e^{-k\theta(t_{i+1}-t_i)}$, and all other transition probabilities are zero under the FM assumption. Because it becomes more likely that multiple events occur as the length of time between observations, dt, increases, we expect probabilities computed under FM to diverge substantially from the Monte Carlo estimates as we increase dt.

We compute Monte Carlo approximations of transition probabilities from 2000 realizations of a BDS process without covariates, with rates $\lambda = 0.0188, \mu = 0.0147, \nu = 0.00268$. These rates are equal to estimates of a transposable element birth, death, and shift rates obtained by Rosenberg et al. [2003] using the FM method. We begin each simulation with an initial population size of 10, and record the state of the process after simulating for dt units of time, varying dt from 0.5 to 10. The approximate transition probability $\hat{p}_{(10,0),(k,l)}(t)$ is then empirically computed by dividing the number of realizations ending in state $\mathbf{X}(t) = (k, l)$ by the total number of simulated processes.

In Figure 2, we see that as the length of an observation interval increases, FM approximations become inaccurate, while those obtained using our method remain within the narrow Monte Carlo confidence intervals. However, notice that the probability that no event occurs remains accurate even under the FM approximation, supporting the efficient implementation of our EM algorithm described in Section 2.5.3. Figure C-1 in the Appendix C demonstrates that our method also reliably calculates other transition probabilities that are set to 0 by the FM method, and these computations remain accurate as we vary the rates of the process.

Further, the discrepancies in numerical transition probabilities between methods indeed translate to differences in estimated rates. To see this, we generate a partially observed dataset and infer rates using both methods. We simulate from the BDS process with parameters $\lambda = 0.07, \mu = 0.12, \nu = 0.02$ to resemble the dynamics of the real dataset we will analyze in the next section, and record 200 discretely observed states of the process evenly spaced dt time units apart. Each simulated interval begins with an initial population size drawn uniformly between 1 and 15, and this data generating process is repeated three times, producing three datasets corresponding to inter-observation intervals of lengths $dt = (0.2, 0.4, 0.6)$. We infer the MLE rates for each of the three discretely observed datasets using the generating function method and under the frequent modeling assumption. This entire procedure is then repeated over 200 trials.

In the top row of Figure 3, we see that our generating function approach successfully recovers the MLE estimates, and coverage of 95% confidence intervals remains close to 0.95 as we increase the length of time intervals between observations. The FM method performs somewhat reasonably for shorter observation intervals, but the bias in these approximate MLEs becomes stark as dt increases, with 95% confidence interval coverage probability dropping as low as 0.24.

Similarly to this transition probability experiment, we check the accuracy of restricted moment computations via simulation by verifying the equality

$$\mathrm{E}(N_t^+ \mid X_0 = i, j) = \sum_{k,l} \mathrm{E}(N_t^+, 1_{x_t=kl} \mid x_0 = i, j)$$

for expected births, and analogous expressions for other expected sufficient statistics. The left hand side is empirically approximated by a Monte Carlo average of the number of births over many realizations of the process, while the restricted moments on the right hand side are the quantities computed via our generating function approach (see Appendix C, Figure C-2).

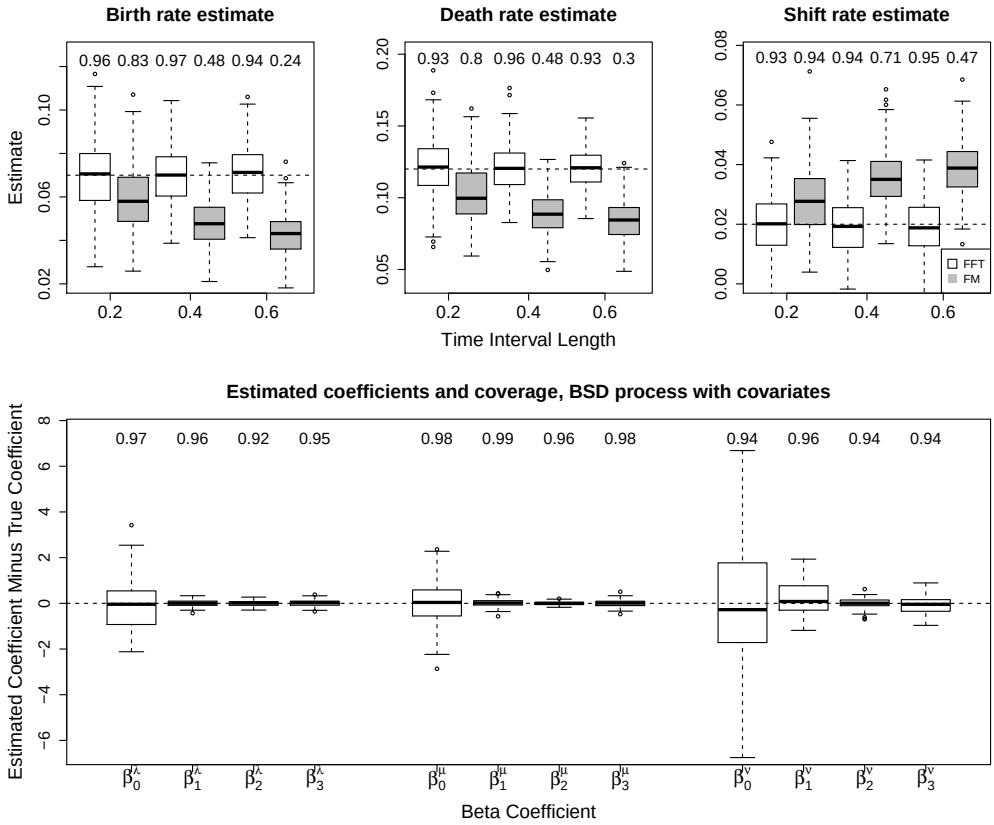

Figure 3: MLE parameter estimates on simulated data. The top row displays estimates of global birth, death, and shift rates in the simple BDS for three datasets, each with observation interval lengths $dt = (0.2, 0.4, 0.6)$. True parameter values used to initialize simulations marked by horizontal dashed line, and results using the FM method are included in gray. Monte Carlo coverage probabilities for 95% confidence intervals are displayed above box plots. The bottom row displays estimated coefficients using EM in the BDS process with covariates, shifted by true values.

2.8 Estimation of parameters in BDS model with covariates

With accurate transition probabilities and restricted moments in hand, we are ready to infer coefficients in the BDS model with covariate-dependent rates using the EM algorithm. We again begin by generating simulated data to resemble the dataset analyzed in the next section. The simulated dataset consists of observations corresponding to 100 "patients", each with three covariates $z_{p,1}, z_{p,2}, z_{p,3} \sim \text{Unif}\{(0,2) \times (6,10) \times (4,6)\}$. We then simulate patient-specific BDS processes, beginning with rates λ_p, ν_p, μ_p log-linearly related to a true vector of coefficients $\boldsymbol{\beta}$ as defined in (2). We collect between 2 and 7 observations per patient, each spaced $dt = 0.4$ apart. Each simulated observation interval begins with an initial number of particles uniformly drawn between 2 and 14. Finally, we set true values of the effect sizes: $\boldsymbol{\beta}^\lambda = [\log(7.5), \log(0.5), \log(0.3), \log(3)]$, $\boldsymbol{\beta}^\nu = [\log(0.5), \log(8), \log(0.5), \log(0.9)]$, $\boldsymbol{\beta}^\mu = [\log(4), \log(0.3), \log(0.8), \log(0.9)]$, chosen so that averaging over patients, the overall birth, shift, and death rates of the process are similar to previous studies [Rosenberg et al., 2003, Doss et al., 2013].

The EM algorithm is initialized with $\boldsymbol{\beta}_0 \sim N(\boldsymbol{\beta}, \text{diag}(0.5\boldsymbol{\beta}))$, and the entire procedure of generating the dataset and inferring rates via EM is repeated 150 times. In the bottom row of Figure 3, we see that the MLEs are again unbiased estimates of the true values, with corresponding confidence interval coverage staying close to 95%.

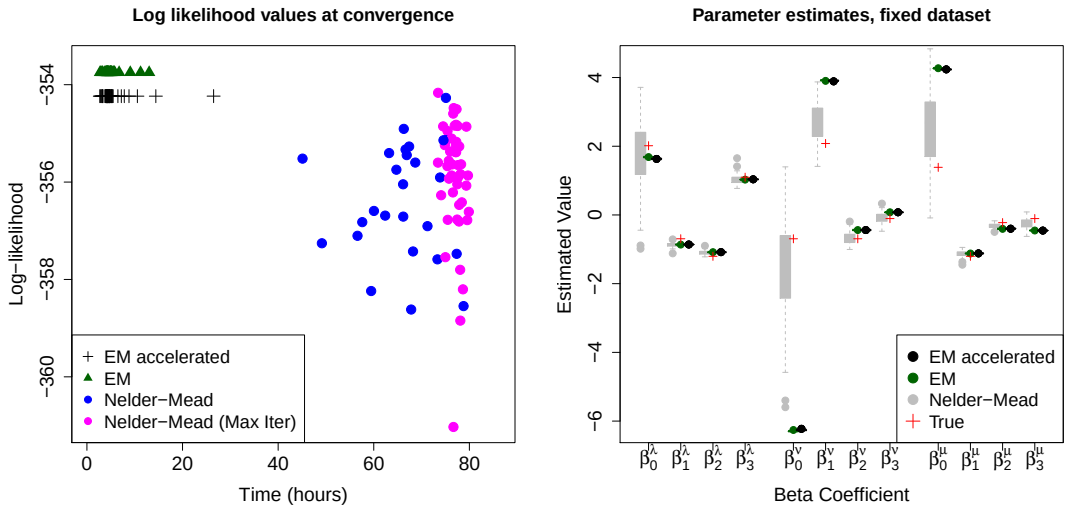

Figure 4: The left plot shows converged log-likelihood values using EM, accelerated EM, and Nelder-Mead optimization. The right plot shows parameter estimates produced by the EM, accelerated EM, and Nelder-Mead algorithms, with true parameters values shown as red crosses.

Having verified that our EM algorithm successfully recovers the true parameters, we turn to a performance comparison with generic optimization via the Nelder-Mead (NM) algorithm implemented in the `optim` package [Nelder and Mead, 1965]. We choose NM as the method for comparison as it proved to be the most robust among the methods available via the `optim` function in R; a similar choice of NM for comparison to EM implementations is motivated in [Lange and Minin, 2013]. In this experiment, we generate one dataset as described in the procedure above from the BDS model with covariates. Fixing these data, we initialize each method with identical initial parameter values and convergence criteria, using a relative tolerance of $\epsilon = 1 \times 10^{-6}$, and repeat this procedure over 100 sets of initial conditions.

Figure 4 displays the log-likelihood values achieved by each algorithm at convergence, as well as values in which Nelder-Mead terminated at an iteration limit set at 2000 steps. We see that in every case, the EM algorithm is significantly faster and finds a better optimum than NM. Further, the wide range of converged log-likelihood values suggests that NM is sensitive to initial conditions — an undesirable feature in this fixed data setting. We also verify that accelerating our EM algorithm according to Section 2.5.3 does not affect the log-likelihood value at convergence up to numerical precision, with a total difference in log-likelihood less than 0.5 accumulated over more than 400 observation intervals. Finally, we note that the comparison between EM and accelerated EM here is included to illustrate that they arrive at the same log-likelihood value and estimates at convergence. The increase in efficiency is not seen here: in these simulated examples, the generating function computations are always performed and cached at each iteration, rather than bypassed for candidate intervals described in Section 2.5.3. In our application to the real dataset in the next section, we find that accelerating EM runs approximately six times as fast as its nonaccelerated counterpart.

Our EM approach is not only more stable in terms of the maximized log-likelihood, but also in terms of parameter estimates. The right panel of Figure 4 shows that estimates for each coefficient differ by no more than 0.01 across disparate initial conditions under both EM implementations, while a range of estimates are produced by the Nelder-Mead algorithm.

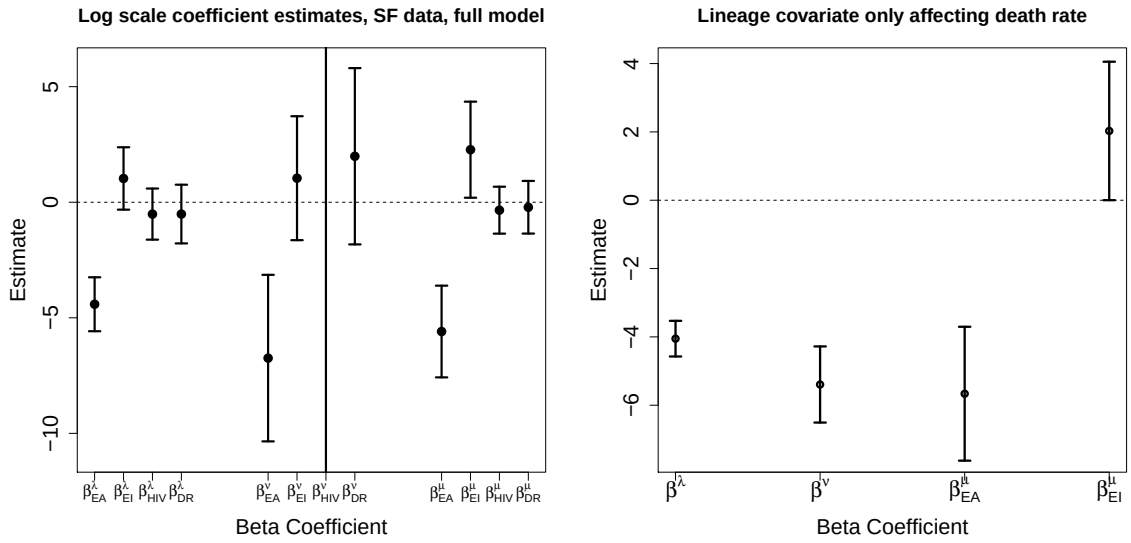

Figure 5: Coefficient estimates and 95% confidence intervals in full model and best model according to BIC. Notice intervals corresponding to β^{μ}_{EI} do not contain 0.

Notice that for some coefficients, estimates produced by NM appear to lie closer to the "true" parameters used to generate the synthetic data. We believe this to be an artifact of centering initial parameter values for both algorithms around the true parameters. Indeed, MLEs corresponding to the likelihood surface of a given *fixed* dataset generally do not exactly coincide with the "true" parameters used to simulate the data. The fact that EM consistently finds a better optimum in terms of log-likelihood demonstrates that this is the case.

2.9 *Mycobacterium tuberculosis* transposable element evolution

We apply our EM algorithm to infer covariate-dependent birth, death, and shift rates of the *M. tuberculosis* transposon IS*6110*, a frequently used marker to track *M. tuberculosis* in the community [McEvoy et al., 2007]. The marker serves as a DNA fingerprint, and in community-based studies patients that share the same or similar *M. tuberculosis* genotypes are considered as part of the same transmission chain [Van Embden et al., 1993, Kato-Maeda et al., 2011]. However, such inference relies on a fairly precise understanding of within-host evolutionary dynamics: for instance, if a DNA marker changes very rapidly, isolates from the same source will be strongly differentiated, and the severity of outbreaks would be underestimated without accounting for the high change rate. Understanding the rates of change of IS*6110*-based genotypes is thus critical toward the interpretation and design of such studies [Tanaka and Rosenberg, 2001], which in turn provide important information toward designing policy decisions such as control and intervention programs.

We analyze data from an ongoing study of the transmission and pathogenesis of *M. tuberculosis* patients in a community study in San Francisco [Cattamanchi et al., 2006, Suwanpimolkul et al., 2013]. The database includes all culture positive tuberculosis cases reported to the San Francisco Department of Public Health. We included patients with more than one *M. tuberculosis* isolate from specimens sampled more than 10 days apart, genotyped with IS*6110* restriction fragment length polymorphism (RFLP) analysis. We assume that changes in the bands marking RFLP patterns evolve according to a linear birth-death-shift process, and assume that patients are not

reinfected with a new strain between observations. Our dataset contains 252 observation intervals corresponding to 196 unique patients observed at 452 time points. Average time between sampling times is 0.35 years, with the longest interval being 2.35 years. Of the 252 intervals, 29 feature end points with distinct genotypes.

This dataset was analyzed by Rosenberg et al. [2003] under the FM assumption, but these authors necessarily discarded all intervals with more than one change in RFLP bands, as these intervals with "complex changes" are not possible under their restricted model. A later investigation by Doss et al. [2013] relaxes this assumption, allowing for multiple births or deaths to occur, but ignores RFLP band locations entirely, working instead only with total copy numbers evolving under a linear birth-death process. Under this birth-death model, the shift rate becomes unidentifiable, and the study instead infers covariate effects of birth and death rates. Our new method allows for a more principled, complete analysis, utilizing the full dataset without compromising any original modeling assumptions.

We begin by applying our EM algorithm to the simple BDS model with a single birth, death, and shift rate of IS*6110* for all patients. We estimate the MLE rates $\hat{\lambda} = 0.0156$, $\hat{\nu} = 0.00426$, $\hat{\mu} = 0.0187$, with associated 95% confidence intervals $(0.00929, 0.0251)$, $(0.00145, 0.0125)$, and $(0.0177, 0.0301)$ respectively. Starting the algorithm from a range of initial parameter values did not affect these results. These estimates are interpretable as the change rate of IS*6110* per copy, per year, and our results are consistent with previous estimates in the literature: for all rates, confidence intervals overlap those obtained in the frequent monitoring approach in [Rosenberg et al., 2003] as well as those obtained in the BD model [Doss et al., 2013]. Similarly to Doss et al. [2013] which estimates $\mu = 0.0207$, we find that our estimate of death rate μ is higher when allowing for multiple events between observations, compared to $\mu = 0.0147$ obtained under the FM assumption. This is to be expected, as there are three intervals in which IS*6110* count drops by more than 1 in the dataset. Although confidence intervals overlap, our estimate of the shift rate is noticeably higher than the previous finding $\nu = 0.00268$ under FM, with the upper end of our confidence interval almost twice as large as the upper end of the 95% FM confidence interval $[0, 0.00654)$. Again, our analysis allows inclusion of several intervals that can be explained by at least two genotype changes that were either omitted in earlier studies or interpreted as a single birth event. Our EM algorithm approach is the first method to our knowledge that is able to accurately estimate the shift rate and produce reliable confidence intervals in the BDS model.

In addition to estimating the BDS rates globally, Doss et al. [2013] investigated rates as functions of several covariates in a panel data setting, and their findings in the birth-death framework suggest that *M. tuberculosis* lineage [Gagneux et al., 2006] may have a statistically significant effect on the rates of the process. We reexamine the effect of lineage on the rates in the full BDS model, considering 109 patients infected with Euro-American (EU) lineage strains, 54 patients with East-Asian (EA) strains, and 25 patients with Indo-Oceanic (IO) strains. We combine EU and IO lineages, because Doss et al. [2013] found that the number of IO samples was not sufficient to recover rates for this lineage. Following Doss et al. [2013], we also include HIV infection status of each patient (HIV) and drug resistance status of the *M. Tuberculosis* strain (DR). These attributes are coded as binary covariates: $\mathrm{EI}_p = 1$ if patient p is infected with the EU or IO strain and 0 otherwise, so that intercept terms $\beta_0^\lambda, \beta_0^\mu, \beta_0^\nu$ correspond to the EA strain. The variable $\mathrm{HIV}_p = 1$ if patient p is infected with HIV and 0 otherwise, and $\mathrm{DR}_p = 1$ if patient p is infected with a drug-resistant strain, and 0 otherwise. Covariates are log-linearly related to birth, death, and shift rates: $\log \lambda_p = \beta_0^\lambda + \beta_1^\lambda \mathrm{EI}_p + \beta_2^\lambda \mathrm{HIV}_p + \beta_3^\lambda \mathrm{DR}_p$, $\log \mu_p = \beta_0^\mu + \beta_1^\mu \mathrm{EI}_p + \beta_2^\mu \mathrm{HIV}_p + \beta_3^\mu \mathrm{DR}_p$, $\log \nu_p = \beta_0^\nu + \beta_1^\nu \mathrm{EI}_p + \beta_2^\nu \mathrm{HIV}_p + \beta_3^\nu \mathrm{DR}_p$.

We estimate coefficients in the full log-linear model described above, as well as in several simpler models, using the EM algorithm. The simpler models differ from the full model by either excluding

Model	# Params	Log-likelihood	BIC
Full, separate EU, IO lineages	15	-119.845	330.01
Full	12	-120.498	313.25
Full, simple ν	9	-122.455	299.10
Lineage covariate only	6	-123.649	293.42
Lineage only, simple ν	5	-123.717	277.54
Lineage only, simple $\boldsymbol{\lambda}, \boldsymbol{\nu}$	**4**	**-124.472**	**273.02**
Simple λ, ν, μ	3	-127.914	273.90

Table 1: Model comparison via BIC $\approx -2\hat{L} + k \ln n$. We also fit the log-linear model in Doss et al. [2013], which includes separate indicator variables for Euro-American and Indo-Oceanic lineages. Models described as "lineage only" do not include HIV, DR covariates, and rates described as "simple" are global to all patients, not influenced by covariates in the model.

the HIV and DR covariates, or excluding all covariates for specified global or "simple" rates. For instance, the model labeled "Lineage only, simple ν" in Table 1 has five parameters $\boldsymbol{\beta} = (\beta_0^\lambda, \beta_1^\lambda, \beta_0^\mu, \beta_1^\mu, \beta^\nu)$, and rates defined as

$$\log \lambda_p = \beta_0^\lambda + \beta_1^\lambda \mathrm{EI}_p, \qquad \log \nu_p = \log \nu = \beta^\nu, \qquad \log \mu_p = \beta_0^\mu + \beta_1^\mu \mathrm{EI}_p.$$

In all cases, estimates obtained using the accelerated EM algorithm and regular implementation coincide, and neither is sensitive to initial conditions. A summary and model comparison via the Bayesian Information Criterion (BIC) [Schwarz, 1978] is included in Table 1, which selects the model including only the lineage covariate for modeling death rate μ. Coefficient estimates are displayed graphically for the full model as well as the best model selected by BIC in Figure 5. While we choose not to report coefficient estimates from each model for brevity, in *all* models, the confidence interval for β_{EI}^μ does not contain zero, indicating that strain lineage has a statistically significant effect on the death rate. The estimate $\hat{\beta_{EI}^\mu} = 2.028$ under the best model indicates that in Euro-American and Indo-Oceanic lineages loss of IS*6110* element occurs $\exp(2.028) = 7.599$ times faster than in their East-Asian counterpart. Our analysis affirms the result suggested by Doss et al. [2013] in the simpler BD framework: *M. tuberculosis* lineage needs to be taken into consideration when studying disease transmission using IS*6110* genotypes.

3 Discussion

In this paper, we have developed an EM algorithm for inference in a discretely observed, multi-type branching process framework. We focus our attention on fitting BDS processes to panel data, driven by the problem of estimating evolutionary dynamics of IS*6110* — a genetic marker that plays an important role in DNA fingerprinting of *M. tuberculosis*. Our method allows for birth, death, and shift log-rates to be linear combinations of many patient-specific covariates, and is flexible enough to capture the full range of dynamics between observation times by approximating the BDS process with a two-type branching process. To our knowledge, there is no other method of comparable accuracy for fitting BDS processes in this setting.

The generating functions we derive and numerical techniques that use these functions to calculate previously unavailable transition probabilities and restricted moments are helpful tools toward probabilistic characterization of such processes more generally. We demonstrate how our generating function approach leads to maximum likelihood estimation and evaluation of expected complete-data log-likelihood within an EM algorithm, but note that these calculations also arise in a variety of

other statistical techniques for prediction and estimation. For example, availability and tractability of the likelihood via our methods allows for their use in Bayesian inference.

Several problems associated with our numerical methods remain open. First, although we have empirical evidence that our branching process approximation to the discretely observed BDS likelihood is very accurate, rigorous characterization of this approximation is lacking. Filling this theoretical gap is an interesting avenue for future research. Second, our method has potential numerical limitations in settings with high population counts. Computing transition probabilities to population sizes up to N of any particle type typically requires N^p differential equations to be solved, where p is the number of particle types. Although efficient numerical solvers are available and each ODE evaluation can be accomplished in only fractions of a second, requiring millions of evaluations becomes prohibitive, especially within an iterative algorithm. However, because the support of transition probabilities is often concentrated unless observation intervals are very long, future work may harness this sparsity to accelerate computations.

We apply our method to analyze within-host evolution of the transposon IS*6110*, an important marker in genetic fingerprinting of *M. tuberculosis*. We obtain confidence intervals for global birth, death, and shift rates λ, μ, ν that overlap with those obtained by Rosenberg et al. [2003] and with those for λ and μ obtained by Doss et al. [2013]. Thus, our estimates are consistent with previous results. While this suggests that the restrictive model assumptions in earlier approaches are not unreasonable in this application, we draw attention to our significantly higher estimate of ν. Indeed, while the frequent monitoring study by Rosenberg et al. [2003] suggests that the global shift rate ν is an order of magnitude smaller than the birth and death rates, our method reveals that the shift rate is in fact comparable to the baseline Euro-American death rate after accounting for strain lineage. This is clearly illustrated in Figure 5, and suggests that the non-negligible shift event should not be omitted from the model as it was in [Doss et al., 2013]. This novel observation was not possible using existing methodology — our approach is the first to accurately estimate the birth, death, and shift rates as functions of covariates in this discretely monitored setting without compromising model assumptions.

Our covariate-specific rate analysis reaffirms previous indication in the simplified BD framework that strain lineage has a significant effect on the death rate [Doss et al., 2013], although the large confidence intervals suggest that this lineage effect is somewhat marginal. Indeed, more data would be required to be certain in the result, but our principled analysis is assuring in that any spurious result can now be attributed to limited, noisy data rather than to model misspecification. The possibility of differences in rates of genetic marker evolution across lineages is important in epidemiological studies. For example, similar IS*6110* genotypes across multiple individuals infected with EA lineage of *M. tuberculosis* do not provide strong evidence of these individuals belonging to the same transmission chain, because of the slow change rate of IS*6110* in the EA lineage. Failing to account for this may lead to inferring false relationships among genotypically similar clusters of patients.

The BDS model we consider is general enough so that our methods can be applied to studying evolution of any transposable element. Such studies are not limited to infectious disease surveillance, because studying evolution of transposable elements in eukaryotes is also of great interest [Biémont, 2010]. Beyond the BDS framework, the tools we develop for fitting branching processes are transferable to many settings. For example, our methodology is applicable to compartmental models, a class of well-known multi-type branching processes that finds applications in modeling cancerous growth, bacterial evolution, and cellular differentiation in systems such as hematopoiesis [Gibson and Renshaw, 1998, Golinelli et al., 2006].

References

G Alsmeyer. On the Galton-Watson predator-prey process. The Annals of Applied Probability, pages 198–211, 1993.

NTJ Bailey. The Elements of Stochastic Processes; with Applications to the Natural Sciences. New York: Wiley, 1964.

NTJ Bailey. The Elements of Stochastic Processes with Applications to the Natural Sciences, volume 25. John Wiley & Sons, 1990.

C Biémont. A brief history of the status of transposable elements: From junk DNA to major players in evolution. Genetics, 186(4):1085–1093, 2010.

JC Butcher. The Numerical Analysis of Ordinary Differential Equations: Runge-Kutta and General Linear Methods. Wiley-Interscience, 1987.

SN Catlin, JL Abkowitz, and P Guttorp. Statistical inference in a two-compartment model for hematopoiesis. Biometrics, 57(2):546–553, 2001.

A Cattamanchi, PC Hopewell, LC Gonzalez, DH Osmond, L Masae Kawamura, CL Daley, and RM Jasmer. A 13-year molecular epidemiological analysis of tuberculosis in San Francisco. The International Journal of Tuberculosis and Lung Disease, 10(3):297–304, 2006.

F Crawford and MA Suchard. Transition probabilities for general birth-death processes with applications in ecology, genetics, and evolution. J. Math. Biol., 65(3):553–580, 2012.

FW Crawford, VN Minin, and MA Suchard. Estimation for general birth-death processes. Journal of the American Statistical Association, 109(506):730–747, 2014.

AP Dempster, NM Laird, and DB Rubin. Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society. Series B (Methodological), 39(1):1–38, 1977.

CR Doss, Ma Suchard, I Holmes, MM Kato-Maeda, and VN Minin. Fitting birth–death processes to panel data with applications to bacterial DNA fingerprinting. The Annals of Applied Statistics, 7(4):2315–2335, 2013.

S Gagneux, K DeRiemer, T Van, M Kato-Maeda, BC de Jong, S Narayanan, M Nicol, S Niemann, K Kremer, MC Gutierrez, et al. Variable host–pathogen compatibility in Mycobacterium tuberculosis. Proceedings of the National Academy of Sciences of the United States of America, 103(8):2869–2873, 2006.

GJ Gibson and E Renshaw. Estimating parameters in stochastic compartmental models using Markov chain methods. Mathematical Medicine and Biology, 15(1):19–40, 1998.

D Golinelli. Bayesian inference in hidden stochastic population processes. PhD thesis, University of Washington, 2000.

D Golinelli, P Guttorp, and JA Abkowitz. Bayesian inference in a hidden stochastic two-compartment model for feline hematopoiesis. Mathematical Medicine and Biology, 23(3):153–172, 2006.

P Guttorp. Stochastic modeling of scientific data. CRC Press, 1995.

P Henrici. Fast Fourier methods in computational complex analysis. Siam Review, 21(4):481–527, 1979.

M Huber. Spatial birth–death swap chains. Bernoulli, 18(3):1031–1041, 2012.

J Illian, A Penttinen, H Stoyan, and D Stoyan. Statistical analysis and modelling of spatial point patterns, volume 70. John Wiley & Sons, 2008.

CH Jackson. Multi-state models for panel data: the msm package for R. Journal of Statistical Software, 38 (8):1–29, 2011.

JD Kalbfleisch and JF Lawless. The analysis of panel data under a Markov assumption. Journal of the American Statistical Association, 80(392):863–871, 1985.

S Karlin and J McGregor. Linear growth birth and death processes. J. Math. Mech., (7):643–662, 1958.

M Kato-Maeda, JZ Metcalfe, and L Flores. Genotyping of *Mycobacterium tuberculosis*: application in epidemiologic studies. Future Microbiology, 6(2):203–216, 2011.

N Keiding. Maximum likelihood estimation in the birth-and-death process. The Annals of Statistics, 3(2): 363–372, 1975.

DG Kendall. On the generalized birth-and-death process. The Annals of Mathematical Statistics, 19(1): 1–15, 1948.

JM Lange and VN Minin. Fitting and interpreting continuous-time latent Markov models for panel data. Statistics in Medicine, 32(26):4581–4595, 2013.

K Lange. Calculation of the equilibrium distribution for a deleterious gene by the finite Fourier transform. Biometrics, 38(1):79–86, 1982.

K Lange. A gradient algorithm locally equivalent to the EM algorithm. Journal of the Royal Statistical Society. Series B. Methodological, 57(2):425–437, 1995.

CRE McEvoy, NCG van Pittius, TC Victor, PD van Helden, and RM Warren. The role of IS*6110* in the evolution of *Mycobacterium tuberculosis*. Tuberculosis, 87(5):393–404, 2007.

VN Minin and MA Suchard. Counting labeled transitions in continuous-time Markov models of evolution. Journal of Mathematical Biology, 56(3):391–412, 2008.

JA Murphy and MR O'donohoe. Some properties of continued fractions with applications in Markov processes. IMA Journal of Applied Mathematics, 16(1):57–71, 1975.

JA Nelder and R Mead. A simplex method for function minimization. The Computer Journal, 7(4):308–313, 1965.

MF Neuts. Algorithmic Probability: A Collection of Problems, volume 3. CRC Press, 1995.

J Neyman, T Park, and EL Scott. Struggle for existence. the Tribolium model: Biological and statistical aspects. In Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, volume 4, pages 41–79. Univ of California Press, 1956.

PS Puri. Interconnected birth and death processes. Journal of Applied Probability, pages 334–349, 1968.

E Renshaw. Stochastic Population Processes: Analysis, Approximations, Simulations. Oxford University Press Oxford, UK, 2011.

NA Rosenberg, AG Tsolaki, and MM Tanaka. Estimating change rates of genetic markers using serial samples: applications to the transposon IS*6110* in *Mycobacterium tuberculosis*. Theoretical Population Biology, 63(4):347–363, 2003.

G Schwarz. Estimating the dimension of a model. The Annals of Statistics, 6(2):461–464, 1978.

G Suwanpimolkul, LG Jarlsberg, JA Grinsdale, D Osmond, LM Kawamura, PC Hopewell, and M Kato-Maeda. Molecular epidemiology of tuberculosis in foreign-born persons living in San Francisco. American Journal of Respiratory and Critical Care Medicine, 187(9):998–1006, 2013.

MM Tanaka and NA Rosenberg. Optimal estimation of transposition rates of insertion sequences for molecular epidemiology. Statistics in Medicine, 20(16):2409–2420, 2001.

JD Van Embden, MD Cave, JT Crawford, JW Dale, KD Eisenach, B Gicquel, P Hermans, C Martin, R McAdam, and TM Shinnick. Strain identification of *Mycobacterium tuberculosis* by DNA fingerprinting: recommendations for a standardized methodology. Journal of Clinical Microbiology, 31(2):406–409, 1993.

Appendix A

Here we derive and solve the Kolmogorov backward equations of the two-type branching process necessary for evaluating the generating functions whose coefficients yield transition probabilities. See [Bailey, 1990] for an exposition on this solution technique.

Our two-type branching process is represent by a vector $(X_1(t), X_2(t))$ that denotes the numbers of particles of two types at time t. Recall the quantities $a_1(k,l)$, the rates of producing k type 1 particles and l type 2 particles, starting with one type 1 particle, and $a_2(k,l)$, analogously defined but beginning with one type 2 particle. Then we may introduce respective pseudo-generating functions $u_i(s_1,s_2) = \sum_k \sum_l a_i(k,l) s_1^k s_2^l$ for $i = 1, 2$, and the probability generating functions can be expressed

$$\begin{aligned}\phi_{10}(t,s_1,s_2) &= E\left[s_1^{X_1(t)} s_2^{X_2(t)} \mid X_1(0)=1, X_2(0)=0\right] = \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} P_{(1,0),(k,l)}(t) s_1^k s_2^l \\ &= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} \left[\mathbf{1}_{k=1,l=0} + a_1(k,l)t + o(t)\right] s_1^k s_2^l = s_1 + u_1(s_1,s_2)t + o(t). \qquad \text{(A-1)}\end{aligned}$$

An analogous expression for $\phi_{01}(t,s_1,s_2)$ is obtained similarly. For short, we write $\phi_{10} := \phi_1, \phi_{01} := \phi_2$, and thus we have the following relations between ϕ and u

$$\frac{d\phi_1(t,s_1,s_2)}{dt}\Big|_{t=0} = u_1(s_1,s_2), \qquad \frac{d\phi_2(t,s_1,s_2)}{dt}\Big|_{t=0} = u_2(s_1,s_2).$$

By particle independence, $\phi_{i,j} = \phi_1^i \phi_2^j$, so it suffices to work with only ϕ_1, ϕ_2. We now derive the backward equations for ϕ_1 and ϕ_2. Chapman-Kolmogorov equations yield the symmetric relations

$$\begin{aligned}\phi_1(t+h,s_1,s_2) &= \phi_1(t,\phi_1(h,s_1,s_2),\phi_2(h,s_1,s_2)) \qquad \text{(A-2)} \\ &= \phi_1(h,\phi_1(t,s_1,s_2),\phi_2(t,s_1,s_2)). \qquad \text{(A-3)}\end{aligned}$$

To derive the backward equations, we begin by expanding $\phi_1(t+h,s_1,s_2)$ around t and applying (A-3):

$$\begin{aligned}\phi_1(t+h,s_1,s_2) &= \phi_1(t,s_1,s_2) + \frac{d\phi_1(t+h,s_1,s_2)}{dh}\Big|_{h=0} h + o(h) \\ &= \phi_1(t,s_1,s_2) + \frac{d\phi_1(h,\phi_1(t,s_1,s_2),\phi_2(t,s_1,s_2))}{dh}\Big|_{h=0} h + o(h) \\ &= \phi_1(t,s_1,s_2) + u_1(\phi_1(t,s_1,s_2),\phi_2(t,s_1,s_2))h + o(h).\end{aligned}$$

Since an analogous argument applies for ϕ_2, we arrive at the system

$$\begin{cases} \frac{d}{dt}\phi_1(t,s_1,s_2) = u_1(\phi_1(t,s_1,s_2),\phi_2(t,s_1,s_2)), \\ \frac{d}{dt}\phi_2(t,s_1,s_2) = u_2(\phi_1(t,s_1,s_2),\phi_2(t,s_1,s_2)), \end{cases}$$

subject to initial conditions $\phi_1(0, s_1, s_2) = s_1, \phi_2(0, s_1, s_2) = s_2$.

We now substitute the rates specific to our birth-shift-death model into this general form: recall the rates defining the two-type branching process formulation presented in Section 2.4 of the main paper are

$$\begin{aligned} &a_1(1,1) = \lambda, \qquad a_1(0,1) = \nu, \qquad a_1(0,0) = \mu \qquad a_1(1,0) = -(\lambda+\nu+\mu), \\ &a_2(0,2) = \lambda, \qquad a_2(0,1) = -(\lambda+\mu), \qquad a_2(0,0) = \mu, \end{aligned} \tag{A-4}$$

so that the pseudo-generating functions and backward equations are

$$\begin{cases} u_1(s_1,s_2) = \lambda s_1 s_2 + \nu s_2 + \mu - (\lambda+\nu+\mu)s_1, & \frac{d}{dt}\phi_1 = \lambda\phi_1\phi_2 + \nu\phi_2 + \mu - (\lambda+\nu+\mu)s_1, \\ u_2(s_1,s_2) = \lambda s_2^2 - (\lambda+\mu)s_2 + \mu, & \frac{d}{dt}\phi_2 = \lambda\phi_2^2 - (\lambda+\mu)\phi_2 + \mu. \end{cases} \tag{A-5}$$

Upon rearranging, the expression for ϕ_2 becomes a Ricatti equation

$$\phi_2' - \lambda\phi_2^2 + (\lambda+\mu)\phi_2 = \mu,$$

and the constant solutions $\phi_2 = 1, \mu/\lambda$ are both particular solutions. Using the simpler root $\phi_2 = 1$, we can reduce the above Ricatti equation to a linear ODE by making a substitution $z = \frac{1}{\phi_2 - 1}$, so that $\phi_2 = 1 + \frac{1}{z}$:

$$\begin{aligned} \phi_2' = -\frac{z'}{z^2} = \mu - (\lambda+\mu)(\frac{1}{z}+1) + \lambda(1+\frac{1}{z})^2 \qquad &= \mu - \frac{\lambda+\mu}{z} - (\lambda+\mu) + \lambda(\frac{1}{z^2} + \frac{2}{z} + 1) \\ = -\frac{\mu-\lambda}{z} + \frac{\lambda}{z^2}. \end{aligned}$$

Multiplying through by $-z^2$ and rearranging, we arrive at a linear equation that is easily solved via the integrating factor method:

$$z' + (\lambda-\mu)z = -\lambda \Rightarrow z = -\frac{\lambda}{\lambda-\mu} + Ce^{-(\lambda-\mu)t}.$$

Substituting ϕ_2 back into the expression, we obtain

$$\phi_2 = 1 + \frac{1}{\frac{\lambda}{\mu-\lambda} + Ce^{(\mu-\lambda)t}},$$

and plugging in the initial condition $\phi_2(0, s_1, s_2) = s_2$, we see $C = \frac{1}{s_2-1} + \frac{\lambda}{\lambda-\mu}$. Thus, we arrive at the closed form solution

$$\phi_2(t, s_1, s_2) = 1 + \left[\frac{\lambda}{\mu-\lambda} + (\frac{1}{s_2-1} + \frac{\lambda}{\lambda-\mu})e^{(\mu-\lambda)t}\right]^{-1} := g(t, s_1, s_2) \tag{A-6}$$

We can now plug this solution into the ODE for ϕ_1 to obtain

$$\frac{d}{dt}\phi_1 + (\lambda+\nu+\mu-\lambda g)\phi_1 = \nu g + \mu. \tag{A-7}$$

Closed form solution for ϕ_1

Equation (A-7) is linear with variable coefficients, and can again be solved by multiplying by an integrating factor. If we define the integrating factor $\psi := \exp\left[\int (\lambda + \nu + \mu - \lambda g)dt\right]$, then

$$\frac{d}{dt}(\phi_1 \psi) = \psi(\nu g + \mu),$$

and after integration and rearranging,

$$\phi_1 = \psi^{-1}\left[\int \psi(\nu g + \mu)dt + C\right]. \tag{A-8}$$

After further simplification, we may write

$$\psi = e^{(\nu+\mu)t}(\lambda s_2 - \mu) + \lambda e^{(\lambda+\nu)t}(1 - s_2),$$

and the integrand becomes

$$\psi(\nu g + \mu) = (\nu + \mu)\psi + \frac{\nu\psi}{\frac{\lambda}{\mu-\lambda} + (\frac{1}{s_2-1} + \frac{\lambda}{\lambda-\mu})e^{(\mu-\lambda)t}}. \tag{A-9}$$

Integrating (A-9) and plugging into (A-8) with initial condition $\phi_1(0, s_1, s_2) = s_1$, we ultimately obtain a closed form expression

$$\begin{aligned}
\phi_1(t, s_1, s_2) = &\left[e^{(\nu+\mu)t}(\lambda s_2 - \mu) + \lambda e^{(\lambda+\nu)t}(1 - s_2)\right]^{-1} \\
&\cdot \Bigg\{\nu(\mu - \lambda)e^{\nu t}\Bigg[\frac{e^{\mu t}(\lambda s_2 - \mu)\,{}_2F_1\Big(1, \frac{\mu+\nu}{\mu-\lambda}, \frac{\lambda-2\mu-\nu}{\lambda-\mu}, \frac{e^{(\mu-\lambda)t}(\lambda s_2-\mu)}{\lambda(s_2-1)}\Big)}{\lambda(\mu+\nu)} \\
&+ \frac{e^{\lambda t}(1 - s_2)\,{}_2F_1\Big(1, \frac{\lambda+\nu}{\mu-\lambda}, \frac{\mu+\nu}{\mu-\lambda}, \frac{e^{(\mu-\lambda)t}(\lambda s_2-\mu)}{\lambda(s_2-1)}\Big)}{\lambda+\nu}\Bigg] + (\lambda s_2 - \mu)e^{(\mu+\nu)t} \\
&+ \frac{\lambda(\nu+\mu)(1-s_2)e^{(\lambda+\nu)t}}{\lambda+\nu} + \mu + s_1(\lambda - \mu) - \lambda s_2 + \frac{\lambda(s_2-1)(\nu+\mu)}{\lambda+\nu} \\
&+ \nu(\lambda - \mu)\Bigg[\frac{\lambda s_2 - \mu}{\lambda(\mu+\nu)}\,{}_2F_1\Big(1, \frac{\mu+\nu}{\mu-\lambda}, \frac{\lambda-2\mu-\nu}{\lambda-\mu}, \frac{\lambda s_2-\mu}{\lambda(s_2-1)}\Big) \\
&+ \frac{1-s_2}{\lambda+\nu}\,{}_2F_1\Big(1, \frac{\lambda+\nu}{\mu-\lambda}, \frac{\mu+\nu}{\mu-\lambda}, \frac{\lambda s_2-\mu}{\lambda(s_2-1)}\Big)\Bigg]\Bigg\},
\end{aligned} \tag{A-10}$$

where ${}_2F_1$ indicates the hypergeometric function. In practice, we solve for ϕ_1 numerically rather than using this closed form solution: evaluating (A-7) via Runge-Kutta methods proves more stable than numerical evaluation of the hypergeometric functions arising in (A-10).

Appendix B

Here we derive the equations in our main theorem. The formulation is repeated below:

Theorem 3.1 *Let $\{X_t\}$ be a two-type branching defined by the rates in equation* (A-4)*. Denote particle time and the number of births, shifts, and deaths over the interval $[0,t)$ by R_t, b_t, f_p, and d_t respectively. Define the generating functions corresponding to births as*

$$H_1^+(r,s_1,s_2,t) = E\left[r^{b_t}s_1^{X_1(t)}s_2^{X_2(t)} \mid \mathbf{X}(0)=(1,0)\right] \text{ and}$$
$$H_2^+(r,s_1,s_2,t) = E\left[r^{b_t}s_1^{X_1(t)}s_2^{X_2(t)} \mid \mathbf{X}(0)=(0,1)\right].$$

Then

$$H_2^+ = y_b + \left[\frac{-\lambda r}{2\lambda r y_b - \lambda - \mu} + \left(\frac{1}{s_2 - y_b} + \frac{\lambda r}{2\lambda r y_b - \lambda - \mu}\right) e^{-(2y_b\lambda r - \lambda - \mu)t}\right]^{-1},$$

where $y_b = (\lambda + \mu + \sqrt{\lambda^2 + 2\lambda\mu + \mu^2 - 4\lambda\mu r})/(2\lambda r)$, and H_1^+ satisfies the following differential equation:

$$\frac{d}{dt}H_1^+(t,s_1,s_2,r) = \lambda r H_1^+ H_2^+ + \nu H_2^+ + \mu - (\lambda + \mu + \nu)H_1^+, \tag{B-1}$$

subject to initial condition $H_1(r,s_1,s_2,0) = s_1$.

The analogous generating functions for shifts, deaths, and particle time satisfy the following equations:

$$H_2^-(t,s_1,s_2,r) = y_d + \left[\frac{-\lambda}{2\lambda y_d - \lambda - \mu} + \left(\frac{1}{s_2 - y_d} + \frac{\lambda}{2\lambda y_d - \lambda - \mu}\right) e^{-(2y_d\lambda - \lambda - \mu)t}\right]^{-1},$$
$$H_2^{\rightarrow}(t,s_1,s_2,r) = 1 + \left[\frac{\lambda}{\mu - \lambda} + (\frac{1}{s_2 - 1} + \frac{\lambda}{\lambda - \mu})e^{(\mu - \lambda)t}\right]^{-1},$$
$$H_2^*(t,s_1,s_2,r) = y_* + \left[\frac{-\lambda}{2\lambda y_* - \lambda - \mu - r} + \left(\frac{1}{s_2 - y_*} + \frac{\lambda}{2\lambda y_* - \lambda - \mu - r}\right) e^{-(2y_*\lambda - \lambda - \mu - r)t}\right]^{-1},$$

$$\frac{d}{dt}H_1^-(t,s_1,s_2,r) = \lambda H_1^- H_2^- + \nu H_2^- + \mu r - (\lambda + \mu + \nu)H_1^-,$$
$$\frac{d}{dt}H_1^{\rightarrow}(t,s_1,s_2,r) = \lambda H_1^{\rightarrow} H_2^{\rightarrow} + \nu r H_2^{\rightarrow} + \mu - (\lambda + \mu + \nu)H_1^{\rightarrow},$$
$$\frac{d}{dt}H_1^*(t,s_1,s_2,r) = \lambda H_1^* H_2^* + \nu H_2^* + \mu - (\lambda + \mu + \nu + r)H_1^*,$$

where $y_d = (\lambda + \mu + \sqrt{\lambda^2 + 2\lambda\mu + \mu^2 - 4\lambda\mu r})/(2\lambda)$, $y_* = (\lambda + \mu + r + \sqrt{(\lambda + \mu + r)^2 - 4\lambda\mu})/(2\lambda)$ and $H_1^-(r,s_1,s_2,0) = H_1^{\rightarrow}(r,s_1,s_2,0) = H_1^*(r,s_1,s_2,0) = s_1$.

Proof Begin by expanding

$$H_{10}^+(t,r,s_1,s_2) = \sum_n \sum_k \sum_l Pr(b_t = n, x_t = (k,l) | x_0 = (1,0)) s_1^k s_2^l r^n.$$

Recall the jump rates of the process in equation (A-4): a_1 correspond to the process beginning with 1 type one particle, and a_2 are jump rates starting with 1 type two particle. We can express the probability terms in H_{10}^+ using the same type of first-order decomposition as in equation (A-1); for instance, in the event of a birth,

$$Pr(b_t = 1, x_t = (1,1) | x_0 = (1,0)) = a_1(1,1) + o(t) = \lambda + o(t)$$

and for other values of $n > 1$,

$$Pr(b_t = n, x_t = (1,1)|x_0 = (1,0)) = o(t).$$

In the case of a shift,

$$Pr(b_t = 0, x_t = (0,1)|x_0 = (1,0)) = a_1(0,1) + o(t) = \nu + o(t)$$

and for other values of $n \neq 0$,

$$Pr(b_t = n, x_t = (0,1)|x_0 = (1,0)) = o(t).$$

We see that the r^n term in the series H_{10}^+ is either $r^1 = r$ if exactly one birth occurs, or $r^0 = 1$ as other powers correspond to more than one event and are absorbed into the $o(t)$ term. Thus,

$$\begin{aligned} H_{10}^+(t,r,s_1,s_2) &= \sum_k \sum_l g_{10,kl}(r,t) s_1^k s_2^l = \sum_n \sum_k \sum_l Pr(b_t = n, x_t = (k,l)|x_0 = (1,0)) s_1^k s_2^l r^n \\ &= s_1 + \lambda s_1 s_2 r + \nu s_2 + \mu - (\lambda + \nu + \mu) s_1 + o(t) := s_1 + u_1^b(s_1,s_2)t + o(t) \end{aligned}$$

with u_1^b denoting the pseudo-generating function, similarly to (A-1). With an analogous derivation for u_2^b, we arrive at the system

$$\begin{cases} u_1^b(s_1,s_2) = \lambda r s_1 s_2 + \nu s_2 + \mu - (\lambda + \nu + \mu) s_1 \\ u_2^b(s_1,s_2) = \lambda r s_2^2 - (\lambda + \mu) s_2 + \mu, \end{cases} \tag{B-2}$$

and since

$$\frac{dH_{10}^+(t,r,s_1,s_2)}{dt}\Big|_{t=0} = u_1^b(s1,s2,r), \qquad \frac{dH_{01}^+(t,r,s_1,s_2)}{dt}\Big|_{t=0} = u_2^b(s1,s2,r),$$

we obtain the backward equations system

$$\begin{cases} \frac{d}{dt} H_{10}^+(t,s_1,s_2,r) = u_1^b(H_{10}^+(t,s_1,s_2,r), H_{01}^+(t,s_1,s_2,r)), \\ \frac{d}{dt} H_{01}^+(t,s_1,s_2,r) = u_2^b(H_{10}^+(t,s_1,s_2,r), H_{01}^+(t,s_1,s_2,r)) \end{cases} \tag{B-3}$$

by the same Chapman-Kolmogorov argument used for transition probabilities, subject to initial conditions $H_{10}(t=0,s_1,s_2,r) = s_1$ and $H_{01}(t=0,s_1,s_2,r) = s_2$. The systems for deaths and shifts are derived analogously beginning with this first-order expansion technique, and are respectively given by

$$\begin{cases} u_1^d(s_1,s_2) = \lambda s_1 s_2 + \nu s_2 + r\mu - (\lambda + \nu + \mu) s_1 \\ u_2^d(s_1,s_2) = \lambda s_2^2 - (\lambda + \mu) s_2 + r\mu, \end{cases} \tag{B-4}$$

$$\begin{cases} u_1^{\rightarrow}(s_1,s_2) = \lambda s_1 s_2 + r\nu s_2 + \mu - (\lambda + \nu + \mu) s_1, \\ u_2^{\rightarrow}(s_1,s_2) = \lambda s_2^2 - (\lambda + \mu) s_2 + \mu. \end{cases} \tag{B-5}$$

To derive the system governing the particle time generating function, recall the quantity $q_{ij,kl}^*(x;t) := Pr(R_t \leq x, X(t) = (k,l)|X(0) = (i,j))$, and consider its Laplace-Stieltjes transform

$$V_{ij,kl}(r;t) = \int_0^\infty e^{-rx} \mathrm{d}q_{ij,kl}^*(x;t). \tag{B-6}$$

The Laplace-Stieltjes transform of such a probability distribution corresponding to a *reward function*, where a_{ij} is the reward accrued per unit time spent in state (i,j), satisfies the forward equation

$$\frac{d}{dt}V_{ij,kl}(r;t) = -a_{ij}rV_{ij,kl}(r;t) + \sum_{m=1}^{K}\sum_{n=1}^{K} Q_{ij,mn}V_{ij,kl}(r;t), \tag{B-7}$$

where $\mathbf{Q}$ is the infinitesimal generator of the Markov chain, with finite or countable number of rows and columns and entries $Q_{ij,kl}$ the instantaneous rates of transitioning from state (i,j) to (k,l), and $Q_{ij,ij} = -\sum_{m,n\neq i,j} Q_{ij,mn}$. Following Neuts [Neuts, 1995], we derive the following integral equation:

$$q^*_{ij,kl}(x,t) = \mathbf{1}_{\{ij=kl\}}\mathbf{1}_{\{x\geq a_{ij}t\}}e^{Q_{ij,ij}t} + \sum_{m,n\neq i,j}\int_0^t e^{Q_{ij,ij}u}Q_{ij,mn}q^*_{mn.kl}(x-a_{ij}u, t-u)\mathrm{d}u.$$

Taking the Laplace transform of both sides and denoting $\widetilde{V}_{ij,kl}(r;t) = \int_0^\infty e^{-rx}q^*_{ij,kl}(x;t)\mathrm{d}x$, we obtain

$$\widetilde{V}_{ij,kl}(r,t) = \mathbf{1}_{\{ij=kl\}}r^{-1}\exp\left[(Q_{ij,ij}-a_{ij})t\right] + \sum_{m,n\neq i,j}\int_0^t e^{Q_{ij,ij}u}Q_{ij,mn}\mathrm{d}u\int_{a_{ij}u}^\infty e^{-rx}q^*_{mn,kl}(x-a_{ij}u;t-u)\mathrm{d}x.$$

Making a change of variables $y = x - a_{ij}u$ in the rightmost integral and multiplying both sides by $\exp\left[-(Q_{ij,ij}-a_{ij})t\right]$ yields

$$\exp\left[-(Q_{ij,ij}-a_{ij})t\right]\widetilde{V}_{ij,kl}(r,t) = \frac{1}{r} + \sum_{m,n\neq i,j}\int_0^t \exp\left[-(Q_{ij,ij}-a_{ij})(t-u)\right]Q_{ij,mn}\widetilde{V}_{mn,kl}(r;t-u).$$

Next, make another substitution $v = t-u$ and simplify after differentiating the above equation with respect to t: we arrive at

$$\frac{\partial}{\partial t}\widetilde{V}_{ij,kl}(r;t) = -a_{ij}r\widetilde{V}_{ij,kl}(r;t) + \sum_{m=1}^{K}\sum_{n=1}^{K} Q_{ij,mn}\widetilde{V}_{mn,kl}(r;t).$$

Equation (B-7) then follows from $V_{ij,kl}(r;t) = s\widetilde{V}_{ij,kl}(r;t)$, with $V_{ij,kl}(t)(r;0) = \mathbf{1}_{\{ij=kl\}}$.

The matrix $\mathbf{V}(r;t) := \{V_{ij,kl}(r;t)\}$ can therefore be written as a matrix exponential

$$\mathbf{V}(r;t) := \exp[\mathbf{Q} - \mathrm{diag}(\mathbf{a})r)t] := \exp(\widetilde{\mathbf{Q}}t), \tag{B-8}$$

where $\mathrm{diag}(\mathbf{a})$ is the diagonal matrix with diagonal entries a_{ij}. In our case, $a_{ij} = 1$, since we are interested in particle time and the "reward" that accumulates per unit of time is that quantity of time itself. Strictly speaking we don't need infinite dimensional matrix algebra here, but we use it to simplify our notation.

Note the similarity of equation (B-8) to the matrix exponential corresponding to transition probabilities $\mathbf{P}(t) = \exp(\mathbf{Q}t)$: thus, the system of backward equations for $V_{ij,kl}$ are almost identical to those for transition probabilities $p_{ij,kl}$. The generators $\widetilde{\mathbf{Q}} \neq \mathbf{Q}$ differ only in diagonal entries: instantaneous rates of no event occurring are augmented by an extra r term $\widetilde{Q}_{ij,ij} = -\sum_{m,n\neq i,j} Q_{ij,mn} - r$. The system of backward equations is thus given by

$$\begin{cases} u^*_1(s_1,s_2) = \lambda s_1 s_2 + \nu s_2 + \mu - (\lambda+\nu+\mu+r)s_1, \\ u^*_2(s_1,s_2) = \lambda s_2^2 + \mu - (\lambda+\mu+r)s_2, \end{cases} \tag{B-9}$$

and as we have seen in the derivation for expected births in equation (B-3), this implies that the generating function

$$H_{10}^*(r, s_1, s_2, t) = \sum_k \sum_l \int_0^\infty e^{-rx} \mathrm{d}q_{ij,kl}^*(x; t) = \sum_k \sum_l V_{10,kl}(r, t) s_1^k s_2^l$$

also satisfies the same system.

Reducing the systems

Each of the four systems for births, shifts, deaths, and particle time can be reduced to a single ODE by first solving the second equation analytically. We demonstrate this in the case of the birth equations (B-3), and abbreviate $H_{10}^+ := H_1, H_{01}^+ := H_2$. Plugging (B-2) into (B-3),

$$\begin{cases} \frac{d}{dt} H_1(t, s_1, s_2, r) = \lambda r H_1 H_2 + \nu H_2 + \mu - (\lambda + \nu + \mu) H_1, \\ \frac{d}{dt} H_2(t, s_1, s_2, r) = \lambda r H_2^2 - (\lambda + \mu) H_2 + \mu. \end{cases}$$

The second equation is a Ricatti equation. To solve it, we first identify a constant solution

$$y_b = \frac{\lambda + \mu + \sqrt{\lambda^2 + 2\lambda\mu + \mu^2 - 4\lambda\mu r}}{2\lambda r}$$

obtained by setting

$$\frac{d}{dt} H_2 = 0 = \lambda r H_2^2 - (\lambda + \mu) H_2 + \mu.$$

Next, perform a change of variables $z = \frac{1}{H_2 - y_b}$ so that $H_2 = y_1 + \frac{1}{z}$, and thus

$$\frac{dz}{dt} + (2y_b\lambda r - \lambda - \mu) z = -\lambda r$$

Using the multiplier method with multiplier $\exp\{(2y_b\lambda r - \lambda - \mu)t\}$, we obtain

$$z = e^{-(2y_b\lambda r - \lambda - \mu)t} \left[\int -\lambda r e^{(2\lambda r y_b - \lambda - \mu)t} dt + C \right] = \frac{-\lambda r}{2\lambda r y_b - \lambda - \mu} + C e^{-(2y_b\lambda r - \lambda - \mu)t}.$$

Thus,

$$H_2 = y_b + \frac{1}{z} = y_b + \left[\frac{-\lambda r}{2\lambda r y_b - \lambda - \mu} + C e^{-(2y_b\lambda r - \lambda - \mu)t} \right]^{-1}$$

and from $H_2(0, r, s_1, s_2) = s_2$, we see $C = \frac{1}{s_2 - y_b} + \frac{\lambda r}{2\lambda r y_b - \lambda - \mu}$. Finally, we arrive at the full solution to the second ODE

$$H_2 := g^b(t, s_1, s_2, r) = y_b + \left[\frac{-\lambda r}{2\lambda r y_b - \lambda - \mu} + \left(\frac{1}{s_2 - y_b} + \frac{\lambda r}{2\lambda r y_b - \lambda - \mu} \right) e^{-(2y_b\lambda r - \lambda - \mu)t} \right]^{-1}.$$

Plugging this solution into the equation for H_1, we have a single ODE that is numerically solvable:

$$\frac{d}{dt} H_1^+(t, s_1, s_2, r) = \lambda r H_1^+ g^b + \nu g^b + \mu - (\lambda + \mu + \nu) H_1^+.$$

An analogous solution beginning with Equations (B-4), (B-5), and (B-9) instead of (B-3) and solving the second Ricatti equation is used to simplify the other equation systems, yielding the results presented in Theorem 3.1.

Appendix C

Here we include additional figures that support, but are not crucial to, illustrating our simulation results.

Figure C-6 displays the transition probabilities $p_{(10,0),(ij)}$ for 25 randomly sampled (i, j) pairs with $0 \leq i, j \leq 32$, calculated by our generating function approach alongside their Monte Carlo estimates and confidence intervals. Monte Carlo estimates are based on 5000 realizations beginning with an initial count of 10 with $dt = 1.0, \lambda = .5, \mu = .45$ and ν ranging from 0.3 to 2.0.

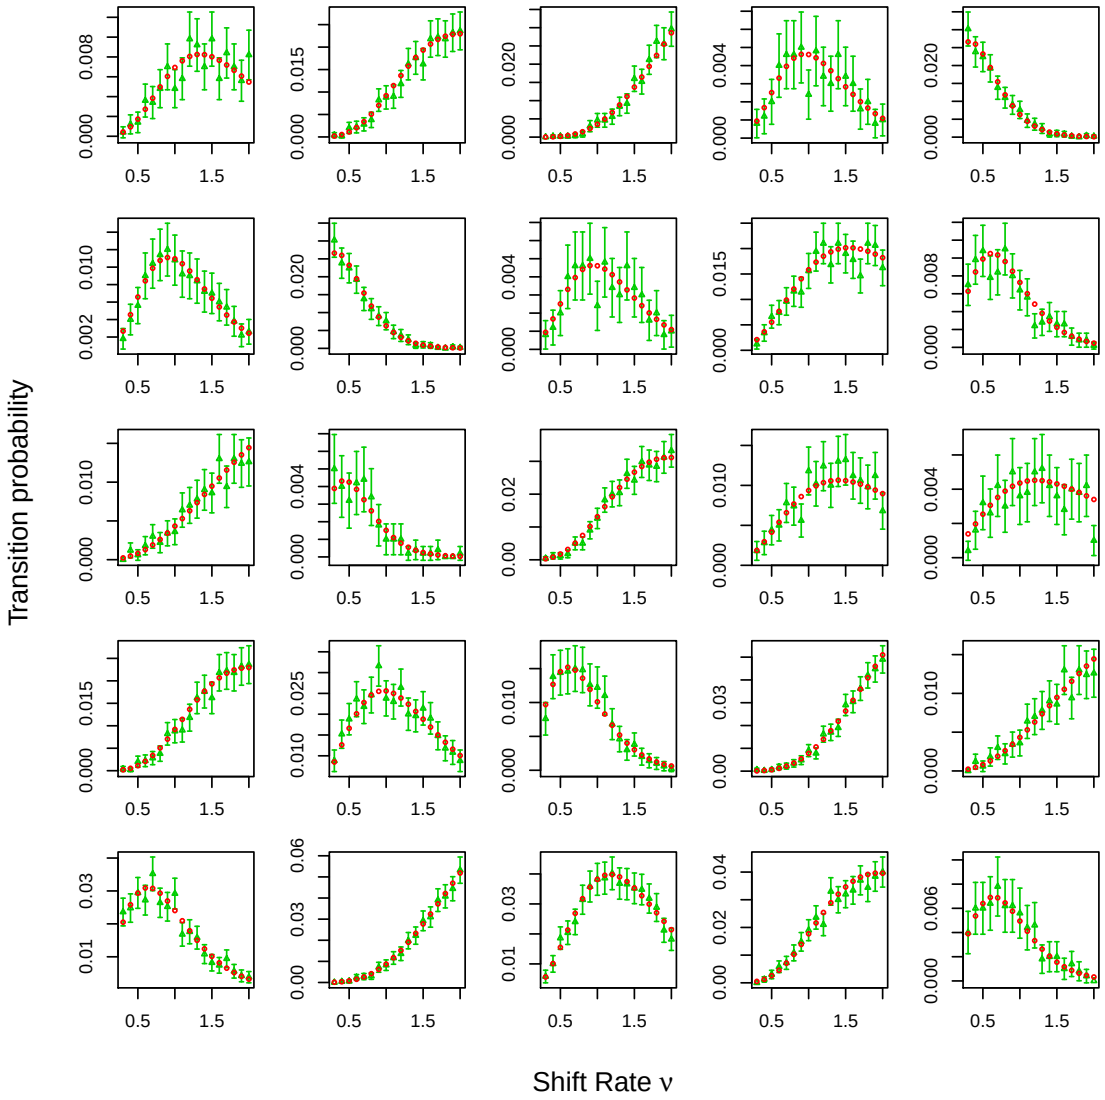

Figure C-6: Transition probabilities remain accurate when increasing all rates of process, presented over a wide range of ν values. Green points and intervals correspond to Monte Carlo estimates of transition probabilities and corresponding 95% confidence intervals. The red points denote probabilities computed with our generating function method.

Figure C-7 shows that restricted moment calculations performed during the E-step are indeed accurate: the following figure corresponds to simulations with 3 times the rates in the Rosenberg-

Tanaka paper: $(\lambda, \nu, \mu) = 3 \cdot (.0188, .0026, .0147)$, with 10 initial particles and varying time lengths.

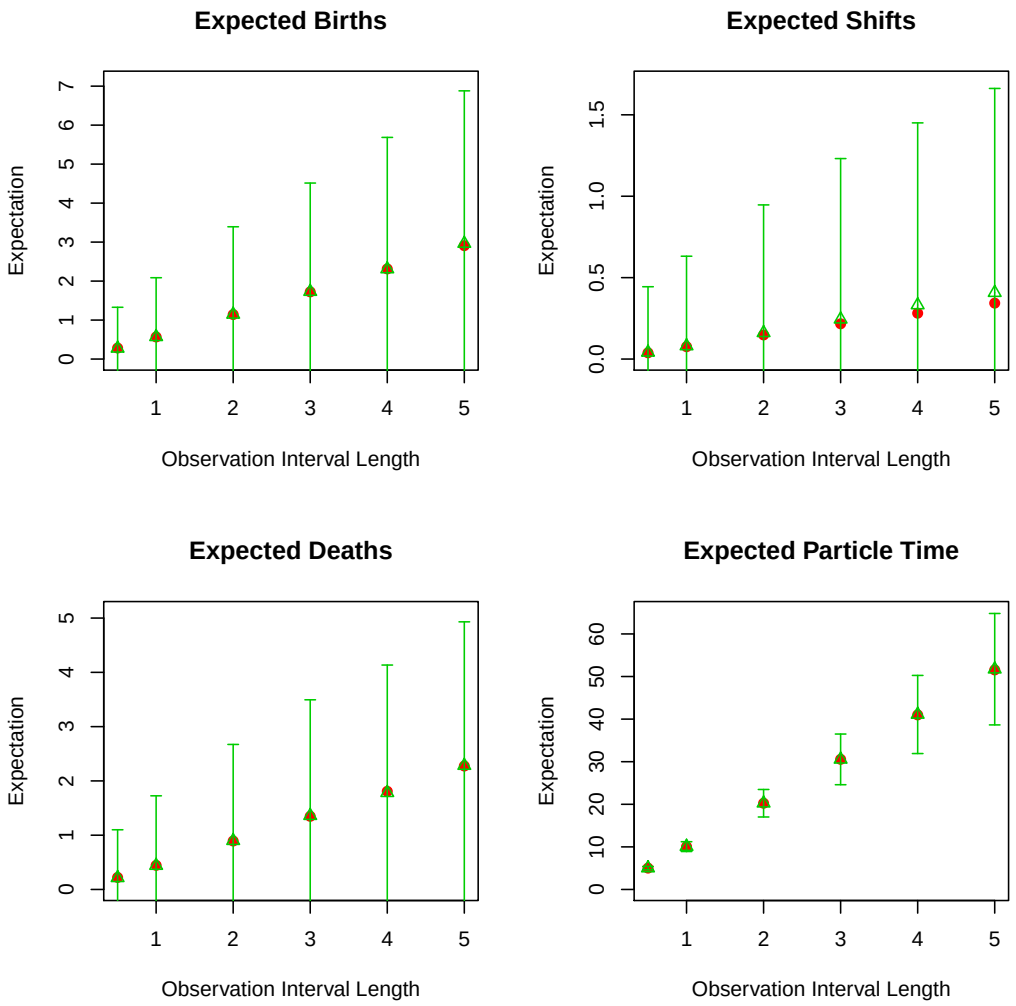

Figure C-7: Restricted moments calculated by our method (red) compared to approximation over 5000 Monte Carlo simulations and corresponding 95% confidence intervals (green).